



OpenHPC (v2.4) Cluster Building Recipes

Rocky 8.5 Base OS

Warewulf/SLURM Edition for Linux* (x86_64)



Document Last Update: 2021-11-16

Document Revision: dc51481d3

Legal Notice

Copyright © 2016-2021, OpenHPC, a Linux Foundation Collaborative Project. All rights reserved.



This documentation is licensed under the Creative Commons Attribution 4.0 International License. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0>.

Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation in the U.S. and/or other countries.
*Other names and brands may be claimed as the property of others.

Contents

| | | |
|----------|--|-----------|
| 1 | Introduction | 5 |
| 1.1 | Target Audience | 5 |
| 1.2 | Requirements/Assumptions | 5 |
| 1.3 | Inputs | 6 |
| 2 | Install Base Operating System (BOS) | 7 |
| 3 | Install OpenHPC Components | 8 |
| 3.1 | Enable OpenHPC repository for local use | 8 |
| 3.2 | Installation template | 9 |
| 3.3 | Add provisioning services on <i>master</i> node | 9 |
| 3.4 | Add resource management services on <i>master</i> node | 10 |
| 3.5 | Optionally add InfiniBand support services on <i>master</i> node | 10 |
| 3.6 | Optionally add Omni-Path support services on <i>master</i> node | 11 |
| 3.7 | Complete basic Warewulf setup for <i>master</i> node | 11 |
| 3.8 | Define <i>compute</i> image for provisioning | 12 |
| 3.8.1 | Build initial BOS image | 12 |
| 3.8.2 | Add OpenHPC components | 12 |
| 3.8.3 | Customize system configuration | 13 |
| 3.8.4 | Additional Customization (<i>optional</i>) | 14 |
| 3.8.4.1 | Enable InfiniBand drivers | 14 |
| 3.8.4.2 | Enable Omni-Path drivers | 14 |
| 3.8.4.3 | Increase locked memory limits | 14 |
| 3.8.4.4 | Enable ssh control via resource manager | 15 |
| 3.8.4.5 | Add BeeGFS | 15 |
| 3.8.4.6 | Add Lustre client | 16 |
| 3.8.4.7 | Enable forwarding of system logs | 16 |
| 3.8.4.8 | Add Nagios monitoring | 17 |
| 3.8.4.9 | Add ClusterShell | 18 |
| 3.8.4.10 | Add <i>genders</i> | 18 |
| 3.8.4.11 | Add Magpie | 18 |
| 3.8.4.12 | Add ConMan | 18 |
| 3.8.4.13 | Add NHC | 19 |
| 3.8.4.14 | Add GEOPM | 19 |
| 3.8.5 | Import files | 20 |
| 3.9 | Finalizing provisioning configuration | 20 |
| 3.9.1 | Assemble bootstrap image | 20 |
| 3.9.2 | Assemble Virtual Node File System (VNFS) image | 20 |
| 3.9.3 | Register nodes for provisioning | 21 |
| 3.9.4 | Optional kernel arguments | 22 |
| 3.9.5 | Optionally configure stateful provisioning | 23 |
| 3.10 | Boot compute nodes | 24 |
| 4 | Install OpenHPC Development Components | 24 |
| 4.1 | Development Tools | 24 |
| 4.2 | Compilers | 25 |
| 4.3 | MPI Stacks | 25 |
| 4.4 | Performance Tools | 26 |
| 4.5 | Setup default development environment | 26 |

| | | |
|-------------------|--|-----------|
| 4.6 | 3rd Party Libraries and Tools | 26 |
| 4.7 | Optional Development Tool Builds | 28 |
| 5 | Resource Manager Startup | 29 |
| 6 | Post-boot compute node configuration | 29 |
| 7 | Run a Test Job | 29 |
| 7.1 | Interactive execution | 30 |
| 7.2 | Batch execution | 32 |
| Appendices | | 33 |
| A | Installation Template | 33 |
| B | Upgrading OpenHPC Packages | 34 |
| C | Integration Test Suite | 35 |
| D | Customization | 37 |
| D.1 | Adding local Lmod modules to OpenHPC hierarchy | 37 |
| D.2 | Rebuilding Packages from Source | 38 |
| E | Package Manifest | 39 |
| F | Package Signatures | 55 |

1 Introduction

This guide presents a simple cluster installation procedure using components from the OpenHPC software stack. OpenHPC represents an aggregation of a number of common ingredients required to deploy and manage an HPC Linux* cluster including provisioning tools, resource management, I/O clients, development tools, and a variety of scientific libraries. These packages have been pre-built with HPC integration in mind while conforming to common Linux distribution standards. The documentation herein is intended to be reasonably generic, but uses the underlying motivation of a small, 4-node stateless cluster installation to define a step-by-step process. Several optional customizations are included and the intent is that these collective instructions can be modified as needed for local site customizations.

Base Linux Edition: this edition of the guide highlights installation without the use of a companion configuration management system and directly uses distro-provided package management tools for component selection. The steps that follow also highlight specific changes to system configuration files that are required as part of the cluster install process.

1.1 Target Audience

This guide is targeted at experienced Linux system administrators for HPC environments. Knowledge of software package management, system networking, and PXE booting is assumed. Command-line input examples are highlighted throughout this guide via the following syntax:

```
[sms]# echo "OpenHPC hello world"
```

Unless specified otherwise, the examples presented are executed with elevated (root) privileges. The examples also presume use of the BASH login shell, though the equivalent commands in other shells can be substituted. In addition to specific command-line instructions called out in this guide, an alternate convention is used to highlight potentially useful tips or optional configuration options. These tips are highlighted via the following format:

Tip

Life is a tale told by an idiot, full of sound and fury signifying nothing. –Willy Shakes

1.2 Requirements/Assumptions

This installation recipe assumes the availability of a single head node *master*, and four *compute* nodes. The *master* node serves as the overall system management server (SMS) and is provisioned with Rocky 8.5 and is subsequently configured to provision the remaining *compute* nodes with Warewulf in a stateless configuration. The terms *master* and SMS are used interchangeably in this guide. For power management, we assume that the compute node baseboard management controllers (BMCs) are available via IPMI from the chosen master host. For file systems, we assume that the chosen master server will host an NFS file system that is made available to the compute nodes. Installation information is also discussed to optionally mount a parallel file system and in this case, the parallel file system is assumed to exist previously.

An outline of the physical architecture discussed is shown in Figure 1 and highlights the high-level networking configuration. The *master* host requires at least two Ethernet interfaces with *eth0* connected to the local data center network and *eth1* used to provision and manage the cluster backend (note that these



Figure 1: Overview of physical cluster architecture.

interface names are examples and may be different depending on local settings and OS conventions). Two logical IP interfaces are expected to each compute node: the first is the standard Ethernet interface that will be used for provisioning and resource management. The second is used to connect to each host's BMC and is used for power management and remote console access. Physical connectivity for these two logical IP networks is often accommodated via separate cabling and switching infrastructure; however, an alternate configuration can also be accommodated via the use of a shared NIC, which runs a packet filter to divert management packets between the host and BMC.

In addition to the IP networking, there is an optional high-speed network (InfiniBand or Omni-Path in this recipe) that is also connected to each of the hosts. This high speed network is used for application message passing and optionally for parallel file system connectivity as well (e.g. to existing Lustre or BeeGFS storage targets).

1.3 Inputs

As this recipe details installing a cluster starting from bare-metal, there is a requirement to define IP addresses and gather hardware MAC addresses in order to support a controlled provisioning process. These values are necessarily unique to the hardware being used, and this document uses variable substitution (`${variable}`) in the command-line examples that follow to highlight where local site inputs are required. A summary of the required and optional variables used throughout this recipe are presented below. Note that while the example definitions above correspond to a small 4-node compute subsystem, the compute parameters are defined in array format to accommodate logical extension to larger node counts.

- `${sms_name}` # Hostname for SMS server
- `${sms_ip}` # Internal IP address on SMS server
- `${sms_eth_internal}` # Internal Ethernet interface on SMS
- `${eth_provision}` # Provisioning interface for computes
- `${internal_netmask}` # Subnet netmask for internal network
- `${ntp_server}` # Local ntp server for time synchronization
- `${bmc_username}` # BMC username for use by IPMI
- `${bmc_password}` # BMC password for use by IPMI
- `${num_computes}` # Total # of desired compute nodes
- `${c_ip[0]}, ${c_ip[1]}, ...` # Desired compute node addresses
- `${c_bmc[0]}, ${c_bmc[1]}, ...` # BMC addresses for computes
- `${c_mac[0]}, ${c_mac[1]}, ...` # MAC addresses for computes
- `${c_name[0]}, ${c_name[1]}, ...` # Host names for computes
- `${compute_regex}` # Regex matching all compute node names (e.g. "c*")
- `${compute_prefix}` # Prefix for compute node names (e.g. "c")

Optional:

- `${sysmgmt_host}` # BeeGFS System Management host name
- `${mgs_fs_name}` # Lustre MGS mount name
- `${sms_ipoib}` # IPoIB address for SMS server
- `${ipoib_netmask}` # Subnet netmask for internal IPoIB
- `${c_ipoib[0]}, ${c_ipoib[1]}, ...` # IPoIB addresses for computes
- `${kargs}` # Kernel boot arguments
- `${nagios_web_password}` # Nagios web access password

2 Install Base Operating System (BOS)

In an external setting, installing the desired BOS on a *master* SMS host typically involves booting from a DVD ISO image on a new server. With this approach, insert the Rocky 8.5 DVD, power cycle the host, and follow the distro provided directions to install the BOS on your chosen *master* host. Alternatively, if choosing to use a pre-installed server, please verify that it is provisioned with the required Rocky 8.5 distribution.

Prior to beginning the installation process of OpenHPC components, several additional considerations are noted here for the SMS host configuration. First, the installation recipe herein assumes that the SMS host name is resolvable locally. Depending on the manner in which you installed the BOS, there may be an adequate entry already defined in `/etc/hosts`. If not, the following addition can be used to identify your SMS host.

```
[sms]# echo ${sms_ip} ${sms_name} >> /etc/hosts
```

While it is theoretically possible to enable SELinux on a cluster provisioned with Warewulf, doing so is beyond the scope of this document. Even the use of permissive mode can be problematic and we therefore recommend disabling SELinux on the *master* SMS host. If SELinux components are installed locally, the `selinuxenabled` command can be used to determine if SELinux is currently enabled. If enabled, consult the distro documentation for information on how to disable.

Finally, provisioning services rely on DHCP, TFTP, and HTTP network protocols. Depending on the local BOS configuration on the SMS host, default firewall rules may prohibit these services. Consequently, this recipe assumes that the local firewall running on the SMS host is disabled. If installed, the default

firewall service can be disabled as follows:

```
[sms]# systemctl disable firewalld
[sms]# systemctl stop firewalld
```

3 Install OpenHPC Components

With the BOS installed and booted, the next step is to add desired OpenHPC packages onto the *master* server in order to provide provisioning and resource management services for the rest of the cluster. The following subsections highlight this process.

3.1 Enable OpenHPC repository for local use

To begin, enable use of the OpenHPC repository by adding it to the local list of available package repositories. Note that this requires network access from your *master* server to the OpenHPC repository, or alternatively, that the OpenHPC repository be mirrored locally. In cases where network external connectivity is available, OpenHPC provides an **ohpc-release** package that includes GPG keys for package signing and enabling the repository. The example which follows illustrates installation of the **ohpc-release** package directly from the OpenHPC build server.

```
[sms]# yum install http://repos.openhpc.community/OpenHPC/2/EL_8/x86_64/ohpc-release-2-1.el8.x86_64.rpm
```

Tip

Many sites may find it useful or necessary to maintain a local copy of the OpenHPC repositories. To facilitate this need, standalone tar archives are provided – one containing a repository of binary packages as well as any available updates, and one containing a repository of source RPMS. The tar files also contain a simple bash script to configure the package manager to use the local repository after download. To use, simply unpack the tarball where you would like to host the local repository and execute the **make_repo.sh** script. Tar files for this release can be found at <http://repos.openhpc.community/dist/2.4>

In addition to the OpenHPC package repository, the *master* host also requires access to the standard base OS distro repositories in order to resolve necessary dependencies. For Rocky 8.5, the requirements are to have access to the BaseOS, Appstream, Extras, PowerTools, and EPEL repositories for which mirrors are freely available online:

- Rocky-8 (e.g. <http://download.rockylinux.org/pub/rocky/8/>)
- EPEL 8 (e.g. <http://download.fedoraproject.org/pub/epel/8/>)

The public EPEL repository will be enabled automatically upon installation of the **ohpc-release** package. Note that this does depend on the Rocky Extras repository, which is shipped with Rocky and is typically enabled by default. In contrast, the PowerTools repository is typically disabled in a standard install, but can be enabled from EPEL as follows:

```
[sms]# yum install dnf-plugins-core
[sms]# yum config-manager --set-enabled powertools
```


3.2 Installation template

The collection of command-line instructions that follow in this guide, when combined with local site inputs, can be used to implement a bare-metal system installation and configuration. The format of these commands is intended to be usable via direct cut and paste (with variable substitution for site-specific settings). Alternatively, the OpenHPC documentation package (`docs-ohpc`) includes a template script which includes a summary of all of the commands used herein. This script can be used in conjunction with a simple text file to define the local site variables defined in the previous section (§ 1.3) and is provided as a convenience for administrators. For additional information on accessing this script, please see Appendix A.

3.3 Add provisioning services on *master* node

With the OpenHPC repository enabled, we can now begin adding desired components onto the *master* server. This repository provides a number of aliases that group logical components together in order to help aid in this process. For reference, a complete list of available group aliases and RPM packages available via OpenHPC are provided in Appendix E. To add support for provisioning services, the following commands illustrate addition of a common base package followed by the Warewulf provisioning system.

```
# Install base meta-packages
[sms]# yum -y install ohpc-base
[sms]# yum -y install ohpc-warewulf
```

Tip

Many server BIOS configurations have PXE network booting configured as the primary option in the boot order by default. If your compute nodes have a different device as the first in the sequence, the `ipmitool` utility can be used to enable PXE.

```
[sms]# ipmitool -E -I lanplus -H ${bmc_ipaddr} -U root chassis bootdev pxe options=persistent
```

HPC systems rely on synchronized clocks throughout the system and the NTP protocol can be used to facilitate this synchronization. To enable NTP services on the SMS host with a specific server `${ntp_server}`, and allow this server to serve as a local time server for the cluster, issue the following:

```
[sms]# systemctl enable chronyd.service
[sms]# echo "server ${ntp_server}" >> /etc/chrony.conf
[sms]# echo "allow all" >> /etc/chrony.conf
[sms]# systemctl restart chronyd
```

Tip

Note that the “allow all” option specified for the chrony time daemon allows all servers on the local network to be able to synchronize with the SMS host. Alternatively, you can restrict access to fixed IP ranges and an example config line allowing access to a local class B subnet is as follows:

```
allow 192.168.0.0/16
```

3.4 Add resource management services on *master* node

OpenHPC provides multiple options for distributed resource management. The following command adds the Slurm workload manager server components to the chosen *master* host. Note that client-side components will be added to the corresponding compute image in a subsequent step.

```
# Install slurm server meta-package
[sms]# yum -y install ohpc-slurm-server

# Use ohpc-provided file for starting SLURM configuration
[sms]# cp /etc/slurm/slurm.conf.ohpc /etc/slurm/slurm.conf

# Identify resource manager hostname on master host
[sms]# perl -pi -e "s/ControlMachine=\S+/ControlMachine=${sms_name}/" /etc/slurm/slurm.conf
```

There are a wide variety of configuration options and plugins available for Slurm and the example config file illustrated above targets a fairly basic installation. In particular, job completion data will be stored in a text file (`/var/log/slurm-jobcomp.log`) that can be used to log simple accounting information. Sites who desire more detailed information, or want to aggregate accounting data from multiple clusters, will likely want to enable the database accounting back-end. This requires a number of additional local modifications (on top of installing `slurm-slurmdbd-ohpc`), and users are advised to consult the online [documentation](#) for more detailed information on setting up a database configuration for Slurm.

Tip

SLURM requires enumeration of the physical hardware characteristics for compute nodes under its control. In particular, three configuration parameters combine to define consumable compute resources: *Sockets*, *CoresPerSocket*, and *ThreadsPerCore*. The default configuration file provided via OpenHPC assumes dual-socket, 8 cores per socket, and two threads per core for this 4-node example. If this does not reflect your local hardware, please update the configuration file at `/etc/slurm/slurm.conf` accordingly to match your particular hardware. Note that the SLURM project provides an easy-to-use online configuration tool that can be accessed [here](#).

Other versions of this guide are available that describe installation of alternate resource management systems, and they can be found in the `docs-ohpc` package.

3.5 Optionally add InfiniBand support services on *master* node

The following command adds OFED and PSM support using base distro-provided drivers to the chosen *master* host.

```
[sms]# yum -y groupinstall "InfiniBand Support"
```

Tip

InfiniBand networks require a subnet management service that can typically be run on either an administrative node, or on the switch itself. The optimal placement and configuration of the subnet manager is beyond the scope of this document, but Rocky 8.5 provides the `opensm` package should you choose to run it on the *master* node.

With the InfiniBand drivers included, you can also enable (optional) IPoIB functionality which provides a mechanism to send IP packets over the IB network. If you plan to mount a Lustre file system over InfiniBand (see §3.8.4.6 for additional details), then having IPoIB enabled is a requirement for the Lustre client. OpenHPC provides a template configuration file to aid in setting up an *ib0* interface on the *master* host. To use, copy the template provided and update the `${sms_ipoib}` and `${ipoib_netmask}` entries to match local desired settings (alter *ib0* naming as appropriate if system contains dual-ported or multiple HCAs).

```
[sms]# cp /opt/ohpc/pub/examples/network/centos/ifcfg-ib0 /etc/sysconfig/network-scripts

# Define local IPoIB address and netmask
[sms]# perl -pi -e "s/master_ipoib/${sms_ipoib}/" /etc/sysconfig/network-scripts/ifcfg-ib0
[sms]# perl -pi -e "s/ipoib_netmask/${ipoib_netmask}/" /etc/sysconfig/network-scripts/ifcfg-ib0

# configure NetworkManager to *not* override local /etc/resolv.conf
[sms]# echo "[main]" > /etc/NetworkManager/conf.d/90-dns-none.conf
[sms]# echo "dns=none" >> /etc/NetworkManager/conf.d/90-dns-none.conf
# Start up NetworkManager to initiate ib0
[sms]# systemctl start NetworkManager
```

3.6 Optionally add Omni-Path support services on *master* node

The following command adds Omni-Path support using base distro-provided drivers to the chosen *master* host.

```
[sms]# yum -y install opa-basic-tools
```

Tip

Omni-Path networks require a subnet management service that can typically be run on either an administrative node, or on the switch itself. The optimal placement and configuration of the subnet manager is beyond the scope of this document, but Rocky 8.5 provides the `opa-fm` package should you choose to run it on the *master* node.

3.7 Complete basic Warewulf setup for *master* node

At this point, all of the packages necessary to use Warewulf on the *master* host should be installed. Next, we need to update several configuration files in order to allow Warewulf to work with Rocky 8.5 and to support local provisioning using a second private interface (refer to Figure 1).

Tip

By default, Warewulf is configured to provision over the `eth1` interface and the steps below include updating this setting to override with a potentially alternatively-named interface specified by `${sms_eth_internal}`.

```
# Configure Warewulf provisioning to use desired internal interface
[sms]# perl -pi -e "s/device = eth1/device = ${sms_eth_internal}/" /etc/warewulf/provision.conf
```

```
# Enable internal interface for provisioning
[sms]# ip link set dev ${sms_eth_internal} up
[sms]# ip address add ${sms_ip}/${internal_netmask} broadcast + dev ${sms_eth_internal}

# Restart/enable relevant services to support provisioning
[sms]# systemctl enable httpd.service
[sms]# systemctl restart httpd
[sms]# systemctl enable dhcpd.service
[sms]# systemctl enable tftp.socket
[sms]# systemctl start tftp.socket
```

3.8 Define *compute* image for provisioning

With the provisioning services enabled, the next step is to define and customize a system image that can subsequently be used to provision one or more *compute* nodes. The following subsections highlight this process.

3.8.1 Build initial BOS image

The OpenHPC build of Warewulf includes specific enhancements enabling support for Rocky 8.5. The following steps illustrate the process to build a minimal, default image for use with Warewulf. We begin by defining a directory structure on the *master* host that will represent the root filesystem of the compute node. The default location for this example is in `/opt/ohpc/admin/images/rocky8.5`.

Tip

Warewulf is configured by default to access an external repository (`download.rockylinux.org`) during the `wwmkchroot` process. If the master host cannot reach the public Rocky mirrors, or if you prefer to access a locally cached mirror, set the `YUM_MIRROR` environment variable to your desired repo location *prior* to running the `wwmkchroot` command below. For example:

```
# Override default OS repository (optional) - set YUM_MIRROR variable to desired repo location
[sms]# export YUM_MIRROR=${BOS_MIRROR}
```

```
# Define chroot location
[sms]# export CHROOT=/opt/ohpc/admin/images/rocky8.5

# Build initial chroot image
[sms]# wwmkchroot -v rocky-8 $CHROOT
# Enable OpenHPC and EPEL repos inside chroot
[sms]# dnf -y --installroot $CHROOT install epel-release
[sms]# cp -p /etc/yum.repos.d/OpenHPC*.repo $CHROOT/etc/yum.repos.d
```

3.8.2 Add OpenHPC components

The `wwmkchroot` process used in the previous step is designed to provide a minimal Rocky 8.5 configuration. Next, we add additional components to include resource management client services, NTP support, and other additional packages to support the default OpenHPC environment. This process augments the chroot-based install performed by `wwmkchroot` to modify the base provisioning image and will access the BOS and OpenHPC repositories to resolve package install requests. We begin by installing a few common base packages:

```
# Install compute node base meta-package
[sms]# yum -y --installroot=$CHROOT install ohpc-base-compute
```

To access the remote repositories by hostname (and not IP addresses), the chroot environment needs to be updated to enable DNS resolution. Assuming that the *master* host has a working DNS configuration in place, the chroot environment can be updated with a copy of the configuration as follows:

```
[sms]# cp -p /etc/resolv.conf $CHROOT/etc/resolv.conf
```

Now, we can include additional required components to the compute instance including resource manager client, NTP, and development environment modules support.

```
# copy credential files into $CHROOT to ensure consistent uid/gids for slurm/munge at
# install. Note that these will be synchronized with future updates via the provisioning system.
[sms]# cp /etc/passwd /etc/group $CHROOT/etc

# Add Slurm client support meta-package and enable munge
[sms]# yum -y --installroot=$CHROOT install ohpc-slurm-client
[sms]# chroot $CHROOT systemctl enable munge

# Register Slurm server with computes (using "configless" option)
[sms]# echo SLURMD_OPTIONS="--conf-server ${sms_ip}" > $CHROOT/etc/sysconfig/slurmd

# Add Network Time Protocol (NTP) support
[sms]# yum -y --installroot=$CHROOT install chrony
# Identify master host as local NTP server
[sms]# echo "server ${sms_ip}" >> $CHROOT/etc/chrony.conf

# Add kernel drivers (matching kernel version on SMS node)
[sms]# yum -y --installroot=$CHROOT install kernel-`uname -r`

# Include modules user environment
[sms]# yum -y --installroot=$CHROOT install lmod-ohpc
```

3.8.3 Customize system configuration

Prior to assembling the image, it is advantageous to perform any additional customization within the chroot environment created for the desired *compute* instance. The following steps document the process to add a local *ssh* key created by Warewulf to support remote access, and enable NFS mounting of a \$HOME file system and the public OpenHPC install path (/opt/ohpc/pub) that will be hosted by the *master* host in this example configuration.

```
# Initialize warewulf database and ssh_keys
[sms]# wwinit database
[sms]# wwinit ssh_keys

# Add NFS client mounts of /home and /opt/ohpc/pub to base image
[sms]# echo "${sms_ip}:/home /home nfs nfsvers=3,nodev,nosuid 0 0" >> $CHROOT/etc/fstab
[sms]# echo "${sms_ip}:/opt/ohpc/pub /opt/ohpc/pub nfs nfsvers=3,nodev 0 0" >> $CHROOT/etc/fstab

# Export /home and OpenHPC public packages from master server
[sms]# echo "/home *(rw,no_subtree_check,fsid=10,no_root_squash)" >> /etc/exports
[sms]# echo "/opt/ohpc/pub *(ro,no_subtree_check,fsid=11)" >> /etc/exports
```

If planning to install the Intel® oneAPI compiler runtime (see §4.7), register the following additional path (`/opt/intel`) to share with computes:

```
# (Optional) Setup NFS mount for /opt/intel if planning to install oneAPI packages
[sms]# mkdir /opt/intel
[sms]# echo "/opt/intel *(ro,no_subtree_check,fsid=12)" >> /etc/exports
[sms]# echo "${sms_ip}:/opt/intel /opt/intel nfs nfsvers=3,nodev 0 0" >> $CHROOT/etc/fstab
```

```
# Finalize NFS config and restart
[sms]# exportfs -a
[sms]# systemctl restart nfs-server
[sms]# systemctl enable nfs-server
```

3.8.4 Additional Customization (*optional*)

This section highlights common additional customizations that can *optionally* be applied to the local cluster environment. These customizations include:

- Add InfiniBand or Omni-Path drivers
- Increase memlock limits
- Restrict ssh access to compute resources
- Add BeeGFS client
- Add Lustre client
- Enable syslog forwarding
- Add Nagios Core monitoring
- Add ClusterShell
- Add *mrsh*
- Add *genders*
- Add ConMan
- Add GEOPM

Details on the steps required for each of these customizations are discussed further in the following sections.

3.8.4.1 Enable InfiniBand drivers If your compute resources support InfiniBand, the following commands add OFED and PSM support using base distro-provided drivers to the compute image.

```
# Add IB support and enable
[sms]# yum -y --installroot=$CHROOT groupinstall "InfiniBand Support"
[sms]# chroot $CHROOT systemctl enable rdma
```

3.8.4.2 Enable Omni-Path drivers If your compute resources support Omni-Path, the following commands add OPA support using base distro-provided drivers to the compute image.

```
# Add OPA support and enable
[sms]# yum -y --installroot=$CHROOT install opa-basic-tools
[sms]# yum -y --installroot=$CHROOT install libpsm2
[sms]# chroot $CHROOT systemctl enable rdma
```

3.8.4.3 Increase locked memory limits In order to utilize InfiniBand or Omni-Path as the underlying high speed interconnect, it is generally necessary to increase the locked memory settings for system users. This can be accomplished by updating the `/etc/security/limits.conf` file and this should be performed within the *compute* image and on all job submission hosts. In this recipe, jobs are submitted from the *master* host, and the following commands can be used to update the maximum locked memory settings on both the master host and the compute image:

```
# Update memlock settings on master
[sms]# perl -pi -e 's/# End of file/\* soft memlock unlimited\n$&/s' /etc/security/limits.conf
[sms]# perl -pi -e 's/# End of file/\* hard memlock unlimited\n$&/s' /etc/security/limits.conf

# Update memlock settings within compute image
[sms]# perl -pi -e 's/# End of file/\* soft memlock unlimited\n$&/s' $CHROOT/etc/security/limits.conf
[sms]# perl -pi -e 's/# End of file/\* hard memlock unlimited\n$&/s' $CHROOT/etc/security/limits.conf
```

3.8.4.4 Enable ssh control via resource manager An additional optional customization that is recommended is to restrict `ssh` access on compute nodes to only allow access by users who have an active job associated with the node. This can be enabled via the use of a pluggable authentication module (PAM) provided as part of the Slurm package installs. To enable this feature within the *compute* image, issue the following:

```
[sms]# echo "account required    pam_slurm.so" >> $CHROOT/etc/pam.d/ssh
```

3.8.4.5 Add BeeGFS To add optional support for mounting BeeGFS file systems, an additional external yum repository provided by the BeeGFS project must be configured. In this recipe, it is assumed that the file system is hosted by servers that are pre-existing and are not part of the install process. The `${sysmgmtd_host}` should point the server running the BeeGFS Management Service. Starting the client service triggers a build of a kernel module, hence the kernel module development packages must be installed first.

```
# Add BeeGFS client software and dependencies to master host
[sms]# wget -P /etc/yum.repos.d https://www.beegfs.io/release/beegfs_7.2.1/dists/beegfs-rhel8.repo
[sms]# yum -y install kernel-devel gcc elfutils-libelf-devel
[sms]# yum -y install beegfs-client beegfs-helperd beegfs-utils

# Enable OFED support in client
[sms]# perl -pi -e "s/^buildArgs=-j8/buildArgs=-j8 BEEGFS_OPENTK_IBVERBS=1/" \
    /etc/beegfs/beegfs-client-autobuild.conf

# Define client's management host
[sms]# /opt/beegfs/sbin/beegfs-setup-client -m ${sysmgmtd_host}

[sms]# systemctl start beegfs-helperd
# Build kernel and mount file system
[sms]# systemctl start beegfs-client
```

```
# Add BeeGFS client software to compute node image
[sms]# wget -P $CHROOT/etc/yum.repos.d https://www.beegfs.io/release/beegfs_7.2.1/dists/beegfs-rhel8.repo
[sms]# yum -y --installroot=$CHROOT install beegfs-client beegfs-helperd beegfs-utils

# Disable auto-build of kernel module in compute node image
[sms]# perl -pi -e "s/^buildEnabled=true/buildEnabled=false/" $CHROOT/etc/beegfs/beegfs-client-autobuild.conf
[sms]# rm -f $CHROOT/var/lib/beegfs/client/force-auto-build

# Enable client daemons on compute nodes
[sms]# chroot $CHROOT systemctl enable beegfs-helperd beegfs-client

# Copy client config to compute nodes
[sms]# cp /etc/beegfs/beegfs-client.conf $CHROOT/etc/beegfs/beegfs-client.conf

# Include kernel module in warewulf bootstrap
[sms]# echo "drivers += beegfs" >> /etc/warewulf/bootstrap.conf
```

3.8.4.6 Add Lustre client To add Lustre client support on the cluster, it is necessary to install the client and associated modules on each host needing to access a Lustre file system. In this recipe, it is assumed that the Lustre file system is hosted by servers that are pre-existing and are not part of the install process. Outlining the variety of Lustre client mounting options is beyond the scope of this document, but the general requirement is to add a mount entry for the desired file system that defines the management server (MGS) and underlying network transport protocol. To add client mounts on both the *master* server and *compute* image, the following commands can be used. Note that the Lustre file system to be mounted is identified by the `${mgs_fs_name}` variable. In this example, the file system is configured to be mounted locally as `/mnt/lustre` .

```
# Add Lustre client software to master host
[sms]# yum -y install lustre-client-ohpc
```

```
# Include Lustre client software in compute image
[sms]# yum -y --installroot=$CHROOT install lustre-client-ohpc

# Include mount point and file system mount in compute image
[sms]# mkdir $CHROOT/mnt/lustre
[sms]# echo "${mgs_fs_name} /mnt/lustre lustre defaults,localflock,noauto,x-systemd.automount 0 0" \
>> $CHROOT/etc/fstab
```

Tip

The suggested mount options shown for Lustre leverage the `localflock` option. This is a [Lustre-specific](#) setting that enables client-local flock support. It is much faster than cluster-wide flock, but if you have an application requiring cluster-wide, coherent file locks, use the standard `flock` attribute instead.

The default underlying network type used by Lustre is *tcp*. If your external Lustre file system is to be mounted using a network type other than *tcp*, additional configuration files are necessary to identify the desired network type. The example below illustrates creation of modprobe configuration files instructing Lustre to use an InfiniBand network with the **o2ib** LNET driver attached to `ib0` . Note that these modifications are made to both the *master* host and *compute* image.

```
[sms]# echo "options lnet networks=o2ib(ib0)" >> /etc/modprobe.d/lustre.conf
[sms]# echo "options lnet networks=o2ib(ib0)" >> $CHROOT/etc/modprobe.d/lustre.conf
```

With the Lustre configuration complete, the client can be mounted on the *master* host as follows:

```
[sms]# mkdir /mnt/lustre
[sms]# mount -t lustre -o localflock ${mgs_fs_name} /mnt/lustre
```

3.8.4.7 Enable forwarding of system logs It is often desirable to consolidate system logging information for the cluster in a central location, both to provide easy access to the data, and to reduce the impact of storing data inside the stateless compute node's memory footprint. The following commands highlight the steps necessary to configure compute nodes to forward their logs to the SMS, and to allow the SMS to accept these log requests.


```
# Configure SMS to receive messages and reload rsyslog configuration
[sms]# echo 'module(load="imudp")' >> /etc/rsyslog.d/ohpc.conf
[sms]# echo 'input(type="imudp" port="514")' >> /etc/rsyslog.d/ohpc.conf
[sms]# systemctl restart rsyslog

# Define compute node forwarding destination
[sms]# echo " *.* @${sms_ip}:514" >> $CHROOT/etc/rsyslog.conf
[sms]# echo "Target=${sms_ip} Protocol=\"udp\" >> $CHROOT/etc/rsyslog.conf

# Disable most local logging on computes. Emergency and boot logs will remain on the compute nodes
[sms]# perl -pi -e "s/^\.*\.info/^\.*\.info/" $CHROOT/etc/rsyslog.conf
[sms]# perl -pi -e "s/^\authpriv/^\authpriv/" $CHROOT/etc/rsyslog.conf
[sms]# perl -pi -e "s/^\mail/^\mail/" $CHROOT/etc/rsyslog.conf
[sms]# perl -pi -e "s/^\cron/^\cron/" $CHROOT/etc/rsyslog.conf
[sms]# perl -pi -e "s/^\uucp/^\uucp/" $CHROOT/etc/rsyslog.conf
```

3.8.4.8 Add Nagios monitoring Nagios is an open source infrastructure network monitoring package designed to watch servers, switches, and various services and offers user-defined alerting facilities for monitoring various aspects of an HPC cluster. The core Nagios daemon and a variety of monitoring plugins are provided by the underlying OS distro and the following commands can be used to install and configure a Nagios server on the *master* node, and add the facility to run tests and gather metrics from provisioned *compute* nodes. This simple configuration example is intended to be illustrative to walk through defining a compute host group and enabling an ssh check for the computes. Users are encouraged to consult Nagios [documentation](#) for more information and can install additional plugins as desired on login nodes, service nodes, or compute hosts.

```
# Install nagios, nrpe, and all available plugins on master host
[sms]# yum -y install --skip-broken nagios nrpe nagios-plugins-*

# Install nrpe and an example plugin into compute node image
[sms]# yum -y --installroot=$CHROOT install nrpe nagios-plugins-ssh

# Enable and configure Nagios NRPE daemon in compute image
[sms]# chroot $CHROOT systemctl enable nrpe
[sms]# perl -pi -e "s/^\allowed_hosts=/# allowed_hosts=/" $CHROOT/etc/nagios/nrpe.cfg
[sms]# echo "nrpe : ${sms_ip} : ALLOW" >> $CHROOT/etc/hosts.allow
[sms]# echo "nrpe : ALL : DENY" >> $CHROOT/etc/hosts.allow

# Copy example Nagios config file to define a compute group and ssh check
# (note: edit as desired to add all desired compute hosts)
[sms]# cp /opt/ohpc/pub/examples/nagios/compute.cfg /etc/nagios/objects
# Register the config file with nagios
[sms]# echo "cfg_file=/etc/nagios/objects/compute.cfg" >> /etc/nagios/nagios.cfg

# Update location of mail binary for alert commands
[sms]# perl -pi -e "s/ \bin/mail/ \usr/bin/mailx/g" /etc/nagios/objects/commands.cfg

# Update email address of contact for alerts
[sms]# perl -pi -e "s/nagios@localhost/root@${sms_name}/" /etc/nagios/objects/contacts.cfg

# Add check_ssh command for remote hosts
[sms]# echo command[check_ssh]=usr/lib64/nagios/plugins/check_ssh localhost $CHROOT/etc/nagios/nrpe.cfg

# define password for nagiosadmin to be able to connect to web interface
[sms]# htpasswd -bc /etc/nagios/passwd nagiosadmin ${nagios_web_password}

# Enable Nagios on master, and configure
[sms]# systemctl enable nagios
```

```
[sms]# systemctl start nagios
# Update permissions on ping command to allow nagios user to execute
[sms]# chmod u+s `which ping`
```

3.8.4.9 Add ClusterShell ClusterShell is an event-based Python library to execute commands in parallel across cluster nodes. Installation and basic configuration defining three node groups (*adm*, *compute*, and *all*) is as follows:

```
# Install ClusterShell
[sms]# yum -y install clustershell

# Setup node definitions
[sms]# cd /etc/clustershell/groups.d
[sms]# mv local.cfg local.cfg.orig
[sms]# echo "adm: ${sms_name}" > local.cfg
[sms]# echo "compute: ${compute_prefix}[1-${num_computes}]" >> local.cfg
[sms]# echo "all: @adm,@compute" >> local.cfg
```

3.8.4.10 Add *genders* *genders* is a static cluster configuration database or node typing database used for cluster configuration management. Other tools and users can access the *genders* database in order to make decisions about where an action, or even what action, is appropriate based on associated types or "genders". Values may also be assigned to and retrieved from a *gender* to provide further granularity. The following example highlights installation and configuration of two genders: *compute* and *bmc*.

```
# Install genders
[sms]# yum -y install genders-ohpc

# Generate a sample genders file
[sms]# echo -e "${sms_name}\tsms" > /etc/genders
[sms]# for ((i=0; i<$num_computes; i++)) ; do
    echo -e "${c_name[$i]}\tcompute,bmc=${c_bmc[$i]}"
done >> /etc/genders
```

3.8.4.11 Add Magpie Magpie contains a number of scripts to aid in running a variety of big data software frameworks within HPC queuing environments. Examples include Hadoop, Spark, Hbase, Storm, Pig, Mahout, Phoenix, Kafka, Zeppelin, and Zookeeper. Consult the online [repository](#) for more information on using these scripts; basic installation is outlined as follows:

```
# Install magpie
[sms]# yum -y install magpie-ohpc
```

3.8.4.12 Add ConMan ConMan is a serial console management program designed to support a large number of console devices and simultaneous users. It supports logging console device output and connecting to compute node consoles via IPMI serial-over-lan. Installation and example configuration is outlined below.

```
# Install conman to provide a front-end to compute consoles and log output
[sms]# yum -y install conman-ohpc

# Configure conman for computes (note your IPMI password is required for console access)
[sms]# for ((i=0; i<$num_computes; i++)) ; do
    echo -n 'CONSOLE name="${c_name[$i]}" dev="ipmi:${c_bmc[$i]}" ' '
```

```

        echo 'ipmiopts="'U:${bmc_username},P:${IPMI_PASSWORD:-undefined},W:so!payloadsize'"
    done >> /etc/conman.conf

# Enable and start conman
[sms]# systemctl enable conman
[sms]# systemctl start conman

```

Note that an additional kernel boot option is typically necessary to enable serial console output. This option is highlighted in §3.9.4 after compute nodes have been registered with the provisioning system.

3.8.4.13 Add NHC Resource managers often provide for a periodic “node health check” to be performed on each compute node to verify that the node is working properly. Nodes which are determined to be “unhealthy” can be marked as down or offline so as to prevent jobs from being scheduled or run on them. This helps increase the reliability and throughput of a cluster by reducing preventable job failures due to misconfiguration, hardware failure, etc. OpenHPC distributes NHC to fulfill this requirement.

In a typical scenario, the NHC driver script is run periodically on each compute node by the resource manager client daemon. It loads its configuration file to determine which checks are to be run on the current node (based on its hostname). Each matching check is run, and if a failure is encountered, NHC will exit with an error message describing the problem. It can also be configured to mark nodes offline so that the scheduler will not assign jobs to bad nodes, reducing the risk of system-induced job failures.

```

# Install NHC on master and compute nodes
[sms]# yum -y install nhc-ohpc
[sms]# yum -y --installroot=${CHROOT} install nhc-ohpc

```

```

# Register as SLURM's health check program
[sms]# echo "HealthCheckProgram=/usr/sbin/nhc" >> /etc/slurm/slurm.conf
[sms]# echo "HealthCheckInterval=300" >> /etc/slurm/slurm.conf # execute every five minutes

```

3.8.4.14 Add GEOPM The Global Extensible Open Power Manager (GEOPM) is a framework for exploring power and energy optimizations targeting high performance computing. The GEOPM package provides built-in features ranging from static management of power policy for each individual compute node, to dynamic coordination of the power policy and performance across all compute nodes hosting an MPI application on a portion of a distributed computing system. The dynamic coordination is implemented as a hierarchical control system for scalable communication and decentralized control. The following commands customize the provisioning environment to support GEOPM installation which is done in a later step in §4.4.

```

# Disable Intel pstate driver for compute nodes as it interferes with GEOPM's operation.
[sms]# export kargs="${kargs} intel_pstate=disable"

```

GEOPM uses the `msr-safe` kernel module to allow users read/write access to whitelisted model specific registers (MSRs). An associated Slurm plugin ensures that MSRs modified within a user’s slurm job are reset to their original state after job completion.

```

# Install msr-safe kernel module and SLURM plugin into compute image
[sms]# yum -y --installroot=${CHROOT} install kmod-msr-safe-ohpc
[sms]# yum -y --installroot=${CHROOT} install msr-safe-ohpc
[sms]# yum -y --installroot=${CHROOT} install msr-safe-slurm-ohpc

```

For documentation on how to configure and use GEOPM, please see the `geopm` man page and tutorials available [online](#).

3.8.5 Import files

The Warewulf system includes functionality to import arbitrary files from the provisioning server for distribution to managed hosts. This is one way to distribute user credentials to *compute* nodes. To import local file-based credentials, issue the following:

```
[sms]# wwsh file import /etc/passwd
[sms]# wwsh file import /etc/group
[sms]# wwsh file import /etc/shadow
```

Similarly, to import the cryptographic key that is required by the *munge* authentication library to be available on every host in the resource management pool, issue the following:

```
[sms]# wwsh file import /etc/munge/munge.key
```

Finally, to add *optional* support for controlling IPoIB interfaces (see §3.5), OpenHPC includes a template file for Warewulf that can optionally be imported and used later to provision *ib0* network settings.

```
[sms]# wwsh file import /opt/ohpc/pub/examples/network/centos/ifcfg-ib0.ww
[sms]# wwsh -y file set ifcfg-ib0.ww --path=/etc/sysconfig/network-scripts/ifcfg-ib0
```

3.9 Finalizing provisioning configuration

Warewulf employs a two-stage boot process for provisioning nodes via creation of a bootstrap image that is used to initialize the process, and a virtual node file system capsule containing the full system image. This section highlights creation of the necessary provisioning images, followed by the registration of desired compute nodes.

3.9.1 Assemble bootstrap image

The bootstrap image includes the runtime kernel and associated modules, as well as some simple scripts to complete the provisioning process. The following commands highlight the inclusion of additional drivers and creation of the bootstrap image based on the running kernel.

```
# (Optional) Include drivers from kernel updates; needed if enabling additional kernel modules on computes
[sms]# export WW_CONF=/etc/warewulf/bootstrap.conf
[sms]# echo "drivers += updates/kernel/" >> $WW_CONF

# Build bootstrap image
[sms]# wwbootstrap `uname -r`
```

3.9.2 Assemble Virtual Node File System (VNFS) image

With the local site customizations in place, the following step uses the `wwvnfs` command to assemble a VNFS capsule from the chroot environment defined for the *compute* instance.

```
[sms]# wwvnfs --chroot $CHROOT
```

3.9.3 Register nodes for provisioning

In preparation for provisioning, we can now define the desired network settings for four example compute nodes with the underlying provisioning system and restart the `dhcp` service. Note the use of variable names for the desired compute hostnames, node IPs, and MAC addresses which should be modified to accommodate local settings and hardware. By default, Warewulf uses network interface names of the `eth#` variety and adds kernel boot arguments to maintain this scheme on newer kernels. Consequently, when specifying the desired provisioning interface via the `$eth_provision` variable, it should follow this convention. Alternatively, if you prefer to use the predictable network interface naming scheme (e.g. names like `en4s0f0`), additional steps are included to alter the default kernel boot arguments and take the `eth#` named interface down after bootstrapping so the normal init process can bring it up again using the desired name.

Also included in these steps are commands to enable Warewulf to manage IPoIB settings and corresponding definitions of IPoIB addresses for the compute nodes. This is typically optional unless you are planning to include a Lustre client mount over InfiniBand. The final step in this process associates the VNFS image assembled in previous steps with the newly defined compute nodes, utilizing the user credential files and munge key that were imported in §3.8.5.

```
# Set provisioning interface as the default networking device
[sms]# echo "GATEWAYDEV=${eth_provision}" > /tmp/network.$$
[sms]# wwsh -y file import /tmp/network.$$ --name network
[sms]# wwsh -y file set network --path /etc/sysconfig/network --mode=0644 --uid=0

# Add nodes to Warewulf data store
[sms]# for ((i=0; i<$num_computes; i++)) ; do
    wwsh -y node new ${c_name[i]} --ipaddr=${c_ip[i]} --hwaddr=${c_mac[i]} -D ${eth_provision}
done
```

```
# Additional step required if desiring to use predictable network interface
# naming schemes (e.g. en4s0f0). Skip if using eth# style names.
[sms]# export kargs="${kargs} net.ifnames=1,biosdevname=1"
[sms]# wwsh provision set --postnetdown=1 "${compute_regex}"
```

```
# Define provisioning image for hosts
[sms]# wwsh -y provision set "${compute_regex}" --vnfs=rocky8.5 --bootstrap='uname -r` \
    --files=dynamic_hosts,passwd,group,shadow,munge.key,network
```

```
# Optionally define IPoIB network settings (required if planning to mount Lustre/BeeGFS over IB)
[sms]# for ((i=0; i<$num_computes; i++)) ; do
    wwsh -y node set ${c_name[i]} -D ib0 --ipaddr=${c_ipoib[i]} --netmask=${ipoib_netmask}
done
[sms]# wwsh -y provision set "${compute_regex}" --fileadd=ifcfg-ib0.ww
```

Tip

Warewulf includes a utility named `wnodescan` to automatically register new compute nodes versus the outlined node-addition approach which requires hardware MAC addresses to be gathered in advance. With `wnodescan`, nodes will be added to the Warewulf database in the order in which their DHCP requests are received by the master, so care must be taken to boot nodes in the order one wishes to see preserved in the Warewulf database. The IP address provided will be incremented after each node is found, and the utility will exit after all specified nodes have been found. Example usage is highlighted below:

```
[sms]# wnodescan --netdev=${eth_provision} --ipaddr=${c_ip[0]} --netmask=${internal_netmask} \
--vnfs=rocky8.5 --bootstrap='uname -r' --listen=${sms_eth_internal} ${c_name[0]}-${c_name[3]}
```

```
# Restart dhcp / update PXE
[sms]# systemctl restart dhcpd
[sms]# wwsh pxe update
```

3.9.4 Optional kernel arguments

The Charliecloud container runtime requires enabling the user namespaces mapping option. This allows applications to run with root privilege inside a container, but have them run as a different, typically non-privileged, user on the host. Though generally regarded as mature and safe, RHEL considers this a Technology Preview, so it must be manually enabled with kernel arguments. We demonstrate enabling on compute nodes, but it would also be required on the SMS if you wish to build and run containers there.

```
# Define node kernel arguments to support user namespaces
[sms]# export kargs="${kargs} namespace.unpriv_enable=1"

# Increase per-user limit on the number of user namespaces that may be created
[sms]# echo "user.max_user_namespaces=15076" >> $CHROOT/etc/sysctl.conf
# rebuild VNFS
[sms]# wvnfs --chroot $CHROOT
```

Tip

Typical Charliecloud workflows are based around Docker containers, but it is not strictly necessary to install Docker itself on the HPC resource. A common pattern is to build the Docker container on a laptop or VM and upload the result to the cluster for use with Charliecloud. More information can be found at <https://hpc.github.io/charliecloud/>

If you chose to enable ConMan in §3.8.4.12, additional warewulf configuration is needed as follows:

```
# Define node kernel arguments to support SOL console
[sms]# wwsh -y provision set "${compute_regex}" --console=ttyS1,115200
```

If any components have added to the boot time kernel command line arguments for the compute nodes, the following command is required to store the configuration in Warewulf:

```
# Set optional compute node kernel command line arguments.
[sms]# wwsh -y provision set "${compute_regex}" --kargs="${kargs}"
```

3.9.5 Optionally configure stateful provisioning

Warewulf normally defaults to running the assembled VNFS image out of system memory in a *stateless* configuration. Alternatively, Warewulf can also be used to partition and format persistent storage such that the VNFS image can be installed locally to disk in a *stateful* manner. This does, however, require that a boot loader (GRUB) be added to the image as follows:

```
# Add GRUB2 bootloader and re-assemble VNFS image
[sms]# yum -y --installroot=$CHROOT install grub2
[sms]# wvvnfs --chroot $CHROOT
```

Enabling stateful nodes also requires additional site-specific, disk-related parameters in the Warewulf configuration, and several example partitioning scripts are provided in the distribution.

```
# Select (and customize) appropriate parted layout example
[sms]# cp /etc/warewulf/filesystem/examples/gpt_example.cmds /etc/warewulf/filesystem/gpt.cmds
[sms]# wwsh provision set --filesystem=gpt "${compute_regex}"
[sms]# wwsh provision set --bootloader=sda "${compute_regex}"
```

Tip

Those provisioning compute nodes in UEFI mode will install a slightly different set of packages in to the VNFS. Warewulf also provides an example EFI filesystem layout.

```
# Add GRUB2 bootloader and re-assemble VNFS image
[sms]# yum -y --installroot=$CHROOT install grub2-efi grub2-efi-modules
[sms]# wvvnfs --chroot $CHROOT
[sms]# cp /etc/warewulf/filesystem/examples/efi_example.cmds /etc/warewulf/filesystem/efi.cmds
[sms]# wwsh provision set --filesystem=efi "${compute_regex}"
[sms]# wwsh provision set --bootloader=sda "${compute_regex}"
```

Upon subsequent reboot of the modified nodes, Warewulf will partition and format the disk to host the desired VNFS image. Once the image is installed to disk, warewulf can be configured to use the nodes' local storage as the boot device.

```
# Configure local boot (after successful provisioning)
[sms]# wwsh provision set --bootlocal=normal "${compute_regex}"
```

3.10 Boot compute nodes

At this point, the *master* server should be able to boot the newly defined compute nodes. Assuming that the compute node BIOS settings are configured to boot over PXE, all that is required to initiate the provisioning process is to power cycle each of the desired hosts using IPMI access. The following commands use the `ipmitool` utility to initiate power resets on each of the four compute hosts. Note that the utility requires that the `IPMI_PASSWORD` environment variable be set with the local BMC password in order to work interactively.

```
[sms]# for ((i=0; i<${num_computes}; i++)) ; do
    ipmitool -E -I lanplus -H ${c_bmc[$i]} -U ${bmc_username} -P ${bmc_password} chassis power reset
done
```

Once kicked off, the boot process should take less than 5 minutes (depending on BIOS post times) and you can verify that the compute hosts are available via `ssh`, or via parallel `ssh` tools to multiple hosts. For example, to run a command on the newly imaged compute hosts using `pdsh`, execute the following:

```
[sms]# pdsh -w c[1-4] uptime
c1 05:03am up 0:02, 0 users, load average: 0.20, 0.13, 0.05
c2 05:03am up 0:02, 0 users, load average: 0.20, 0.14, 0.06
c3 05:03am up 0:02, 0 users, load average: 0.19, 0.15, 0.06
c4 05:03am up 0:02, 0 users, load average: 0.15, 0.12, 0.05
```

Tip

While the `pxelinux.0` and `lpxelinux.0` files that ship with Warewulf to enable network boot support a wide range of hardware, some hosts may boot more reliably or faster using the BOS versions provided via the `syslinux-tftpboot` package. If you encounter PXE issues, consider replacing the `pxelinux.0` and `lpxelinux.0` files supplied with `warewulf-provision-ohpc` with versions from `syslinux-tftpboot`.

4 Install OpenHPC Development Components

The install procedure outlined in §3 highlighted the steps necessary to install a *master* host, assemble and customize a *compute* image, and provision several compute hosts from bare-metal. With these steps completed, additional OpenHPC-provided packages can now be added to support a flexible HPC development environment including development tools, C/C++/FORTRAN compilers, MPI stacks, and a variety of 3rd party libraries. The following subsections highlight the additional software installation procedures.

4.1 Development Tools

To aid in general development efforts, OpenHPC provides recent versions of the GNU autotools collection, the Valgrind memory debugger, EasyBuild, and Spack. These can be installed as follows:

```
# Install autotools meta-package
[sms]# yum -y install ohpc-autotools

[sms]# yum -y install EasyBuild-ohpc
[sms]# yum -y install hwloc-ohpc
[sms]# yum -y install spack-ohpc
[sms]# yum -y install valgrind-ohpc
```


4.2 Compilers

OpenHPC presently packages the GNU compiler toolchain integrated with the underlying Lmod modules system in a hierarchical fashion. The modules system will conditionally present compiler-dependent software based on the toolchain currently loaded.

```
[sms]# yum -y install gnu9-compilers-ohpc
```

4.3 MPI Stacks

For MPI development and runtime support, OpenHPC provides pre-packaged builds for a variety of MPI families and transport layers. Currently available options and their applicability to various network transports are summarized in Table 1. The command that follows installs a starting set of MPI families that are compatible with both ethernet and high-speed fabrics.

Table 1: Available MPI variants

| | Ethernet (TCP) | InfiniBand | Intel® Omni-Path |
|-------------------|----------------|------------|------------------|
| MPICH (ofi) | ✓ | ✓ | ✓ |
| MPICH (ucx) | ✓ | ✓ | ✓ |
| MVAPICH2 | | ✓ | |
| MVAPICH2 (psm2) | | | ✓ |
| OpenMPI (ofi/ucx) | ✓ | ✓ | ✓ |

```
[sms]# yum -y install openmpi4-gnu9-ohpc mpich-ofi-gnu9-ohpc
```

Note that OpenHPC 2.x introduces the use of two related transport layers for the MPICH and OpenMPI builds that support a variety of underlying fabrics: [UCX](#) (Unified Communication X) and [OFI](#) (OpenFabrics interfaces). In the case of OpenMPI, a monolithic build is provided which supports both transports and end-users can customize their runtime preferences with environment variables. For MPICH, two separate builds are provided and the example above highlighted installing the ofi variant. However, the packaging is designed such that both versions can be installed simultaneously and users can switch between the two via normal module command semantics. Alternatively, a site can choose to install the ucx variant instead as a drop-in MPICH replacement:

```
[sms]# yum -y install mpich-ucx-gnu9-ohpc
```

In the case where both MPICH variants are installed, two modules will be visible in the end-user environment and an example of this configuration is highlighted is below.

```
[sms]# module avail mpich
----- /opt/ohpc/pub/moduledeps/gnu9 -----
mpich/3.3.2-ofi  mpich/3.3.2-ucx (D)
```

If your system includes InfiniBand and you enabled underlying support in §3.5 and §3.8.4, an additional MVAPICH2 family is available for use:

```
[sms]# yum -y install mvapich2-gnu9-ohpc
```

Alternatively, if your system includes Intel® Omni-Path, use the (psm2) variant of MVAPICH2 instead:

```
[sms]# yum -y install mvapich2-psm2-gnu9-ohpc
```

4.4 Performance Tools

OpenHPC provides a variety of open-source tools to aid in application performance analysis (refer to Appendix E for a listing of available packages). This group of tools can be installed as follows:

```
# Install perf-tools meta-package
[sms]# yum -y install ohpc-gnu9-perf-tools
```

Optionally, the GEOPM power management framework can be installed using the convenience meta-package below. Note that GEOPM requires customization of the compute node environment to include an additional kernel module as highlighted previously in §3.8.4.14):

```
# Install GEOPM meta-package
[sms]# yum -y install ohpc-gnu9-geopm
```

4.5 Setup default development environment

System users often find it convenient to have a default development environment in place so that compilation can be performed directly for parallel programs requiring MPI. This setup can be conveniently enabled via modules and the OpenHPC modules environment is pre-configured to load an `ohpc` module on login (if present). The following package install provides a default environment that enables autotools, the GNU compiler toolchain, and the OpenMPI stack.

```
[sms]# yum -y install lmod-defaults-gnu9-openmpi4-ohpc
```

Tip

If you want to change the default environment from the suggestion above, OpenHPC also provides the GNU compiler toolchain with the MPICH and MVAPICH2 stacks:

- `lmod-defaults-gnu9-mpich-ofi-ohpc`
- `lmod-defaults-gnu9-mpich-ucx-ohpc`
- `lmod-defaults-gnu9-mvapich2-ohpc`

4.6 3rd Party Libraries and Tools

OpenHPC provides pre-packaged builds for a number of popular open-source tools and libraries used by HPC applications and developers. For example, OpenHPC provides builds for FFTW and HDF5 (including serial and parallel I/O support), and the GNU Scientific Library (GSL). Again, multiple builds of each package are available in the OpenHPC repository to support multiple compiler and MPI family combinations where appropriate. Note, however, that not all combinatorial permutations may be available for components where there are known license incompatibilities. The general naming convention for builds provided by OpenHPC is to append the compiler and MPI family name that the library was built against directly into the package

name. For example, libraries that do not require MPI as part of the build process adopt the following RPM name:

```
package-<compiler_family>-ohpc-<package_version>-<release>.rpm
```

Packages that do require MPI as part of the build expand upon this convention to additionally include the MPI family name as follows:

```
package-<compiler_family>-<mpi_family>-ohpc-<package_version>-<release>.rpm
```

To illustrate this further, the command below queries the locally configured repositories to identify all of the available PETSc packages that were built with the GNU toolchain. The resulting output that is included shows that pre-built versions are available for each of the supported MPI families presented in §4.3.

```
[sms]# yum search petsc-gnu9 ohpc
Loaded plugins: fastestmirror
Loading mirror speeds from cached hostfile
===== N/S matched: petsc-gnu9, ohpc =====
petsc-gnu9-impi-ohpc.x86_64 : Portable Extensible Toolkit for Scientific Computation
petsc-gnu9-mpich-ohpc.x86_64 : Portable Extensible Toolkit for Scientific Computation
petsc-gnu9-mvapich2-ohpc.x86_64 : Portable Extensible Toolkit for Scientific Computation
petsc-gnu9-openmpi4-ohpc.x86_64 : Portable Extensible Toolkit for Scientific Computation
```

Tip

OpenHPC-provided 3rd party builds are configured to be installed into a common top-level repository so that they can be easily exported to desired hosts within the cluster. This common top-level path (`/opt/ohpc/pub`) was previously configured to be mounted on *compute* nodes in §3.8.3, so the packages will be immediately available for use on the cluster after installation on the *master* host.

For convenience, OpenHPC provides package aliases for these 3rd party libraries and utilities that can be used to install available libraries for use with the GNU compiler family toolchain. For parallel libraries, aliases are grouped by MPI family toolchain so that administrators can choose a subset should they favor a particular MPI stack. Please refer to Appendix E for a more detailed listing of all available packages in each of these functional areas. To install all available package offerings within OpenHPC, issue the following:

```
# Install 3rd party libraries/tools meta-packages built with GNU toolchain
[sms]# yum -y install ohpc-gnu9-serial-libs
[sms]# yum -y install ohpc-gnu9-io-libs
[sms]# yum -y install ohpc-gnu9-python-libs
[sms]# yum -y install ohpc-gnu9-runtimes
```

```
# Install parallel lib meta-packages for all available MPI toolchains
[sms]# yum -y install ohpc-gnu9-mpich-parallel-libs
[sms]# yum -y install ohpc-gnu9-openmpi4-parallel-libs
```

4.7 Optional Development Tool Builds

In addition to the 3rd party development libraries built using the open source toolchains mentioned in §4.6, OpenHPC also provides *optional* compatible builds for use with the compilers and MPI stack included in newer versions of the Intel® oneAPI HPC Toolkit (using the *classic* compiler variants). These packages provide a similar hierarchical user environment experience as other compiler and MPI families present in OpenHPC.

To take advantage of the available builds, OpenHPC provides a convenience package to enable the oneAPI repository locally along with compatibility packages that integrate oneAPI-generated compiler and MPI modulefiles within the standard OpenHPC user environment. To enable the Intel® oneAPI repository and install minimum compiler and MPI requirements for OpenHPC packaging, issue the following:

```
# Enable Intel oneAPI and install OpenHPC compatibility packages
[sms]# yum -y install intel-oneapi-toolkit-release-ohpc
[sms]# rpm --import https://yum.repos.intel.com/intel-gpg-keys/GPG-PUB-KEY-INTEL-SW-PRODUCTS.PUB
[sms]# yum -y install intel-compilers-devel-ohpc
[sms]# yum -y install intel-mpi-devel-ohpc
```

Tip

As noted in §3.8.3, the default installation path for OpenHPC (`/opt/ohpc/pub`) is exported over NFS from the *master* to the compute nodes, but the Intel® oneAPI HPC Toolkit packages install to a top-level path of `/opt/intel`. To make the Intel® compilers available to the compute nodes one must add an additional NFS export for `/opt/intel` that is mounted on desired compute nodes.

To enable all 3rd party builds available in OpenHPC that are compatible with the Intel® oneAPI classic compiler suite, issue the following:

```
# Optionally, choose the Omni-Path enabled build for MVAPICH2. Otherwise, skip to retain IB variant
[sms]# yum -y install mvapich2-psm2-intel-ohpc
```

```
# Install 3rd party libraries/tools meta-packages built with Intel toolchain
[sms]# yum -y install ohpc-intel-serial-libs
[sms]# yum -y install ohpc-intel-geopm
[sms]# yum -y install ohpc-intel-io-libs
[sms]# yum -y install ohpc-intel-perf-tools
[sms]# yum -y install ohpc-intel-python3-libs
[sms]# yum -y install ohpc-intel-mpich-parallel-libs
[sms]# yum -y install ohpc-intel-mvapich2-parallel-libs
[sms]# yum -y install ohpc-intel-openmpi4-parallel-libs
[sms]# yum -y install ohpc-intel-impi-parallel-libs
```

5 Resource Manager Startup

In section §3, the Slurm resource manager was installed and configured for use on both the *master* host and *compute* node instances. With the cluster nodes up and functional, we can now startup the resource manager services in preparation for running user jobs. Generally, this is a two-step process that requires starting up the controller daemons on the *master* host and the client daemons on each of the *compute* hosts. Note that Slurm leverages the use of the *munge* library to provide authentication services and this daemon also needs to be running on all hosts within the resource management pool. The following commands can be used to startup the necessary services to support resource management under Slurm.

```
# Start munge and slurm controller on master host
[sms]# systemctl enable munge
[sms]# systemctl enable slurmctld
[sms]# systemctl start munge
[sms]# systemctl start slurmctld

# Start slurm clients on compute hosts
[sms]# pdsh -w $compute_prefix[1-4] systemctl start munge
[sms]# pdsh -w $compute_prefix[1-4] systemctl start slurmd
```

6 Post-boot compute node configuration

```
# Generate NHC configuration file based on compute node environment
[sms]# pdsh -w c1 "/usr/sbin/nhc-genconf -H '*' -c -" | dshbak -c
```

7 Run a Test Job

With the resource manager enabled for production usage, users should now be able to run jobs. To demonstrate this, we will add a “test” user on the *master* host that can be used to run an example job.

```
[sms]# useradd -m test
```

Warewulf installs a utility on the compute nodes to automatically synchronize known files from the provisioning server at five minute intervals. In this recipe, recall that we previously registered credential files with Warewulf (e.g. `passwd`, `group`, and `shadow`) so that these files would be propagated during compute node imaging. However, with the addition of a new “test” user above, the files have been outdated and we need to update the Warewulf database to incorporate the additions. This re-sync process can be accomplished as follows:

```
[sms]# wwsh file resync passwd shadow group
```

Tip

After re-syncing to notify Warewulf of file modifications made on the *master* host, it should take approximately 5 minutes for the changes to propagate. However, you can also manually pull the changes from compute nodes via the following:

```
[sms]# pdsh -w $compute_prefix[1-4] /warewulf/bin/wvgetfiles
```

OpenHPC includes a simple “hello-world” MPI application in the `/opt/ohpc/pub/examples` directory that can be used for this quick compilation and execution. OpenHPC also provides a companion job-launch utility named `prun` that is installed in concert with the pre-packaged MPI toolchains. This convenience script provides a mechanism to abstract job launch across different resource managers and MPI stacks such that a single launch command can be used for parallel job launch in a variety of OpenHPC environments. It also provides a centralizing mechanism for administrators to customize desired environment settings for their users.

7.1 Interactive execution

To use the newly created “test” account to compile and execute the application *interactively* through the resource manager, execute the following (note the use of `prun` for parallel job launch which summarizes the underlying native job launch mechanism being used):

```
# Switch to "test" user
[sms]# su - test

# Compile MPI "hello world" example
[test@sms ~]$ mpicc -O3 /opt/ohpc/pub/examples/mpi/hello.c

# Submit interactive job request and use prun to launch executable
[test@sms ~]$ salloc -n 8 -N 2

[test@c1 ~]$ prun ./a.out

[prun] Master compute host = c1
[prun] Resource manager = slurm
[prun] Launch cmd = mpiexec.hydra -bootstrap slurm ./a.out

Hello, world (8 procs total)
--> Process # 0 of 8 is alive. -> c1
--> Process # 4 of 8 is alive. -> c2
--> Process # 1 of 8 is alive. -> c1
--> Process # 5 of 8 is alive. -> c2
--> Process # 2 of 8 is alive. -> c1
--> Process # 6 of 8 is alive. -> c2
--> Process # 3 of 8 is alive. -> c1
--> Process # 7 of 8 is alive. -> c2
```

Tip

The following table provides approximate command equivalences between SLURM and OpenPBS:

| Command | OpenPBS | SLURM |
|----------------------------------|----------------------|------------------------------|
| Submit <i>batch</i> job | qsub [job script] | sbatch [job script] |
| Request <i>interactive</i> shell | qsub -I /bin/bash | salloc |
| Delete job | qdel [job id] | scancel [job id] |
| Queue status | qstat -q | sinfo |
| Job status | qstat -f [job id] | scontrol show job [job id] |
| Node status | pbsnodes [node name] | scontrol show node [node id] |

7.2 Batch execution

For batch execution, OpenHPC provides a simple job script for reference (also housed in the `/opt/ohpc/pub/examples` directory). This example script can be used as a starting point for submitting batch jobs to the resource manager and the example below illustrates use of the script to submit a batch job for execution using the same executable referenced in the previous interactive example.

```
# Copy example job script
[test@sms ~]$ cp /opt/ohpc/pub/examples/slurm/job.mpi .

# Examine contents (and edit to set desired job sizing characteristics)
[test@sms ~]$ cat job.mpi
#!/bin/bash

#SBATCH -J test                # Job name
#SBATCH -o job.%j.out         # Name of stdout output file (%j expands to %jobId)
#SBATCH -N 2                  # Total number of nodes requested
#SBATCH -n 16                  # Total number of mpi tasks #requested
#SBATCH -t 01:30:00           # Run time (hh:mm:ss) - 1.5 hours

# Launch MPI-based executable

prun ./a.out

# Submit job for batch execution
[test@sms ~]$ sbatch job.mpi
Submitted batch job 339
```

Tip

The use of the `%j` option in the example batch job script shown is a convenient way to track application output on an individual job basis. The `%j` token is replaced with the Slurm job allocation number once assigned (job #339 in this example).

Appendices

A Installation Template

This appendix highlights the availability of a companion installation script that is included with OpenHPC documentation. This script, when combined with local site inputs, can be used to implement a starting recipe for bare-metal system installation and configuration. This template script is used during validation efforts to test cluster installations and is provided as a convenience for administrators as a starting point for potential site customization.

Tip

Note that the template script provided is intended for use during initial installation and is not designed for repeated execution. If modifications are required after using the script initially, we recommend running the relevant subset of commands interactively.

The template script relies on the use of a simple text file to define local site variables that were outlined in §1.3. By default, the template installation script attempts to use local variable settings sourced from the `/opt/ohpc/pub/doc/recipes/vanilla/input.local` file, however, this choice can be overridden by the use of the `${OHPC_INPUT_LOCAL}` environment variable. The template install script is intended for execution on the SMS *master* host and is installed as part of the `docs-ohpc` package into `/opt/ohpc/pub/doc/recipes/vanilla/recipe.sh`. After enabling the OpenHPC repository and reviewing the guide for additional information on the intent of the commands, the general starting approach for using this template is as follows:

1. Install the `docs-ohpc` package

```
[sms]# yum -y install docs-ohpc
```

2. Copy the provided template input file to use as a starting point to define local site settings:

```
[sms]# cp /opt/ohpc/pub/doc/recipes/rocky8/input.local input.local
```

3. Update `input.local` with desired settings
4. Copy the template installation script which contains command-line instructions culled from this guide.

```
[sms]# cp -p /opt/ohpc/pub/doc/recipes/rocky8/x86_64/warewulf/slurm/recipe.sh .
```

5. Review and edit `recipe.sh` to suite.
6. Use environment variable to define local input file and execute `recipe.sh` to perform a local installation.

```
[sms]# export OHPC_INPUT_LOCAL=./input.local
[sms]# ./recipe.sh
```

B Upgrading OpenHPC Packages

As newer OpenHPC releases are made available, users are encouraged to upgrade their locally installed packages against the latest repository versions to obtain access to bug fixes and newer component versions. This can be accomplished with the underlying package manager as OpenHPC packaging maintains versioning state across releases. Also, package builds available from the OpenHPC repositories have “-ohpc” appended to their names so that wild cards can be used as a simple way to obtain updates. The following general procedure highlights a method for upgrading existing installations. When upgrading from a minor release older than v2, you will first need to update your local OpenHPC repository configuration to point against the v2 release (or update your locally hosted mirror). Refer to §3.1 for more details on enabling the latest repository. In contrast, when upgrading between micro releases on the same branch (e.g. from v2 to 2.2), there is no need to adjust local package manager configurations when using the public repository as rolling updates are pre-configured.

1. (Optional) Ensure repo metadata is current (on head node and in chroot location(s)). Package managers will naturally do this on their own over time, but if you are wanting to access updates immediately after a new release, the following can be used to sync to the latest.

```
[sms]# yum clean expire-cache
[sms]# yum --installroot=$CHROOT clean expire-cache
```

2. Upgrade master (SMS) node

```
[sms]# yum -y upgrade "*-ohpc"

# Any new Base OS provided dependencies can be installed by
# updating the ohpc-base metapackage
[sms]# yum -y upgrade "ohpc-base"
```

Tip

The version of Warewulf included in OpenHPC v1.3.6 added a new architecture-specific package containing iPXE files. Upgraders will need to install this package on the SMS node.

```
[sms]# yum -y install warewulf-provision-server-ipxe-x86_64-ohpc
```

3. Upgrade packages in compute image

```
[sms]# yum -y --installroot=$CHROOT upgrade "*-ohpc"

# Any new compute-node Base OS provided dependencies can be installed by
# updating the ohpc-base-compute metapackage
[sms]# yum -y --installroot=$CHROOT upgrade "ohpc-base-compute"
```

4. Rebuild image(s)

```
[sms]# wwnfs --chroot $CHROOT
```

In the case where packages were upgraded within the chroot compute image, you will need to reboot the compute nodes when convenient to enable the changes.

C Integration Test Suite

This appendix details the installation and basic use of the integration test suite used to support OpenHPC releases. This suite is not intended to replace the validation performed by component development teams, but is instead, devised to confirm component builds are functional and interoperable within the modular OpenHPC environment. The test suite is generally organized by components and the OpenHPC CI workflow relies on running the full suite using [Jenkins](#) to test multiple OS configurations and installation recipes. To facilitate customization and running of the test suite locally, we provide these tests in a standalone RPM.

```
[sms]# yum -y install test-suite-ohpc
```

The RPM installation creates a user named `ohpc-test` to house the test suite and provide an isolated environment for execution. Configuration of the test suite is done using standard GNU autotools semantics and the [BATS](#) shell-testing framework is used to execute and log a number of individual unit tests. Some tests require privileged execution, so a different combination of tests will be enabled depending on which user executes the top-level `configure` script. Non-privileged tests requiring execution on one or more compute nodes are submitted as jobs through the SLURM resource manager. The tests are further divided into “short” and “long” run categories. The short run configuration is a subset of approximately 180 tests to demonstrate basic functionality of key components (e.g. MPI stacks) and should complete in 10-20 minutes. The long run (around 1000 tests) is comprehensive and can take an hour or more to complete.

Most components can be tested individually, but a default configuration is setup to enable collective testing. To test an isolated component, use the `configure` option to disable all tests, then re-enable the desired test to run. The `--help` option to `configure` will display all possible tests. By default, the test suite will endeavor to run tests for multiple MPI stacks where applicable. To restrict tests to only a subset of MPI families, use the `--with-mpi-families` option (e.g. `--with-mpi-families="openmpi4"`). Example output is shown below (some output is omitted for the sake of brevity).

```
[sms]# su - ohpc-test
[test@sms ~]$ cd tests
[test@sms ~]$ ./configure --disable-all --enable-fftw
checking for a BSD-compatible install... /bin/install -c
checking whether build environment is sane... yes
...
----- SUMMARY -----
Package version..... : test-suite-2.0.0

Build user..... : ohpc-test
Build host..... : sms001
Configure date..... : 2020-10-05 08:22
Build architecture..... : x86_64
Compiler Families..... : gnu9
MPI Families..... : mpich mvapich2 openmpi4
Python Families..... : python3
Resource manager ..... : SLURM
Test suite configuration..... : short
...
Libraries:
  Adios ..... : disabled
  Boost ..... : disabled
  Boost MPI..... : disabled
  FFTW..... : enabled
  GSL..... : disabled
  HDF5..... : disabled
  HYPRE..... : disabled
...
```

Many OpenHPC components exist in multiple flavors to support multiple compiler and MPI runtime permutations, and the test suite takes this in to account by iterating through these combinations by default. If `make check` is executed from the top-level test directory, all configured compiler and MPI permutations of a library will be exercised. The following highlights the execution of the FFTW related tests that were enabled in the previous step.

```
[test@sms ~]$ make check
make --no-print-directory check-TESTS
PASS: libs/fftw/ohpc-tests/test_mpi_families
=====
Testsuite summary for test-suite 2.0.0
=====
# TOTAL: 1
# PASS: 1
# SKIP: 0
# XFAIL: 0
# FAIL: 0
# XPASS: 0
# ERROR: 0
=====
[test@sms ~]$ cat libs/fftw/tests/family-gnu*/rm_execution.log
1..3
ok 1 [libs/FFTW] Serial C binary runs under resource manager (SLURM/gnu9/mpich)
ok 2 [libs/FFTW] MPI C binary runs under resource manager (SLURM/gnu9/mpich)
ok 3 [libs/FFTW] Serial Fortran binary runs under resource manager (SLURM/gnu9/mpich)
PASS rm_execution (exit status: 0)
1..3
ok 1 [libs/FFTW] Serial C binary runs under resource manager (SLURM/gnu9/mvapich2)
ok 2 [libs/FFTW] MPI C binary runs under resource manager (SLURM/gnu9/mvapich2)
ok 3 [libs/FFTW] Serial Fortran binary runs under resource manager (SLURM/gnu9/mvapich2)
PASS rm_execution (exit status: 0)
1..3
ok 1 [libs/FFTW] Serial C binary runs under resource manager (SLURM/gnu9/openmpi4)
ok 2 [libs/FFTW] MPI C binary runs under resource manager (SLURM/gnu9/openmpi4)
ok 3 [libs/FFTW] Serial Fortran binary runs under resource manager (SLURM/gnu9/openmpi4)
PASS rm_execution (exit status: 0)
```

D Customization

D.1 Adding local Lmod modules to OpenHPC hierarchy

Locally installed applications can easily be integrated in to OpenHPC systems by following the Lmod convention laid out by the provided packages. Two sample module files are included in the `examples-ohpc` package—one representing an application with no compiler or MPI runtime dependencies, and one dependent on OpenMPI and the GNU toolchain. Simply copy these files to the prescribed locations, and the `lmod` application should pick them up automatically.

```
[sms]# mkdir /opt/ohpc/pub/modulefiles/example1
[sms]# cp /opt/ohpc/pub/examples/example.modulefile \
/opt/ohpc/pub/modulefiles/example1/1.0
[sms]# mkdir /opt/ohpc/pub/moduledeps/gnu7-openmpi3/example2
[sms]# cp /opt/ohpc/pub/examples/example-mpi-dependent.modulefile \
/opt/ohpc/pub/moduledeps/gnu7-openmpi3/example2/1.0
[sms]# module avail
```

```
----- /opt/ohpc/pub/moduledeps/gnu7-openmpi3 -----
adios/1.12.0  imb/2018.0      netcdf-fortran/4.4.4  ptscotch/6.0.4    sionlib/1.7.1
boost/1.65.1  mpi4py/2.0.0    netcdf/4.4.1.1       scalapack/2.0.2   slepc/3.7.4
example2/1.0  mpiP/3.4.1      petsc/3.7.6          scalasca/2.3.1    superlu_dist/4.2
fftw/3.3.6    mumps/5.1.1     phdf5/1.10.1         scipy/0.19.1      tau/2.26.1
hypre/2.11.2  netcdf-cxx/4.3.0  pnetcdf/1.8.1        scorep/3.1         trilinos/12.10.1

----- /opt/ohpc/pub/moduledeps/gnu7 -----
R/3.4.2      metis/5.1.0     ocr/1.0.1            pdttoolkit/3.24   superlu/5.2.1
gsl/2.4      mpich/3.2       openblas/0.2.20      plasma/2.8.0
hdf5/1.10.1  numpy/1.13.1    openmpi3/3.0.0 (L)   scotch/6.0.4

----- /opt/ohpc/admin/modulefiles -----
spack/0.10.0

----- /opt/ohpc/pub/modulefiles -----
EasyBuild/3.4.1  cmake/3.9.2    hwloc/1.11.8         pmix/1.2.3         valgrind/3.13.0
autotools        (L)  example1/1.0 (L)  llvm5/5.0.0         prun/1.2           (L)
clustershell/1.8  gnu7/7.2.0 (L)  ohpc                 (L)  singularity/2.4

Where:
L: Module is loaded

Use "module spider" to find all possible modules.
Use "module keyword key1 key2 ..." to search for all possible modules matching any of the "keys".
```

D.2 Rebuilding Packages from Source

Users of OpenHPC may find it desirable to rebuild one of the supplied packages to apply build customizations or satisfy local requirements. One way to accomplish this is to install the appropriate source RPM, modify the spec file as needed, and rebuild to obtain an updated binary RPM. OpenHPC spec files contain macros to facilitate local customizations of compiler, compilation flags and MPI family. A brief example using the FFTW library is highlighted below. Note that the source RPMs can be downloaded from the community repository server at <http://repos.openhpc.community> via a web browser or directly via yum as highlighted below. In this example we make an explicit change to FFTW's configuration, as well as modifying the CFLAGS environment variable. The package is also tagged with an additional delimiter to allow easy co-installation and use.

```
# Install rpm-build package and yum tools from base OS distro
[test@sms ~]$ sudo yum -y install rpm-build yum-utils

# Install FFTW's build dependencies
[test@sms ~]$ sudo yum-builddep fftw-gnu9-openmpi4-ohpc

# Download SRPM from OpenHPC repository and install locally
[test@sms ~]$ yumdownloader --source fftw-gnu9-openmpi4-ohpc
[test@sms ~]$ rpm -i ./fftw-gnu9-openmpi4-ohpc-3.3.8-5.1.ohpc.2.0.src.rpm

# Modify spec file as desired
[test@sms ~]$ cd ~/rpmbuild/SPECS
[test@sms ~rpmbuild/SPECS]$ perl -pi -e "s/enable-static=no/enable-static=yes/" fftw.spec

# Increment RPM release so the package manager will see an update
[test@sms ~rpmbuild/SPECS]$ perl -pi -e "s/Release: 5.1/Release: 6.1/" fftw.spec

# Rebuild binary RPM. Note that additional directives can be specified to modify build
[test@sms ~rpmbuild/SPECS]$ rpmbuild -bb --define "OHPC_CFLAGS '-O3 -mtune=native'" \
    --define "OHPC_CUSTOM_DELIM static" fftw.spec

# Install the new package
[test@sms ~rpmbuild/SPECS]$ sudo yum -y install \
    ../RPMS/x86_64/fftw-gnu9-openmpi4-static-ohpc.2.0-3.3.8-6.1.x86_64.rpm

# The new module file appears along side the default
[test@sms ~]$ module avail fftw

-----/opt/ohpc/pub/moduledeps/gnu9-openmpi4 -----
fftw/3.3.8-static  fftw/3.3.8 (D)
```

E Package Manifest

This appendix provides a summary of available meta-package groupings and all of the individual RPM packages that are available as part of this OpenHPC release. The meta-packages provide a mechanism to group related collections of RPMs by functionality and provide a convenience mechanism for installation. A list of the available meta-packages and a brief description is presented in [Table 2](#).

Table 2: Available OpenHPC Meta-packages

| Group Name | Description |
|----------------------------------|---|
| ohpc-autotools | Collection of GNU autotools packages. |
| ohpc-base | Collection of base packages. |
| ohpc-base-compute | Collection of compute node base packages. |
| ohpc-gnu9-geopm | Global Extensible Open Power Manager for use with GNU compiler toolchain. |
| ohpc-gnu9-io-libs | Collection of IO library builds for use with GNU compiler toolchain. |
| ohpc-gnu9-mpich-io-libs | Collection of IO library builds for use with GNU compiler toolchain and the MPICH runtime. |
| ohpc-gnu9-mpich-parallel-libs | Collection of parallel library builds for use with GNU compiler toolchain and the MPICH runtime. |
| ohpc-gnu9-mpich-perf-tools | Collection of performance tool builds for use with GNU compiler toolchain and the MPICH runtime. |
| ohpc-gnu9-mvapich2-io-libs | Collection of IO library builds for use with GNU compiler toolchain and the MVAPICH2 runtime. |
| ohpc-gnu9-mvapich2-parallel-libs | Collection of parallel library builds for use with GNU compiler toolchain and the MVAPICH2 runtime. |
| ohpc-gnu9-mvapich2-perf-tools | Collection of performance tool builds for use with GNU compiler toolchain and the MVAPICH2 runtime. |
| ohpc-gnu9-openmpi4-io-libs | Collection of IO library builds for use with GNU compiler toolchain and the OpenMPI runtime. |
| ohpc-gnu9-openmpi4-parallel-libs | Collection of parallel library builds for use with GNU compiler toolchain and the OpenMPI runtime. |
| ohpc-gnu9-openmpi4-perf-tools | Collection of performance tool builds for use with GNU compiler toolchain and the OpenMPI runtime. |
| ohpc-gnu9-parallel-libs | Collection of parallel library builds for use with GNU compiler toolchain. |
| ohpc-gnu9-perf-tools | Collection of performance tool builds for use with GNU compiler toolchain. |
| ohpc-gnu9-python-libs | Collection of python related library builds for use with GNU compiler toolchain. |
| ohpc-gnu9-python3-libs | Collection of python3 related library builds for use with GNU compiler toolchain. |
| ohpc-gnu9-runtimes | Collection of runtimes for use with GNU compiler toolchain. |
| ohpc-gnu9-serial-libs | Collection of serial library builds for use with GNU compiler toolchain. |
| ohpc-intel-geopm | Global Extensible Open Power Manager for use with Intel(R) Parallel Studio XE software suite. |
| ohpc-intel-impi-io-libs | Collection of IO library builds for use with Intel(R) Parallel Studio XE software suite and Intel(R) MPI runtime. |
| ohpc-intel-impi-parallel-libs | Collection of parallel library builds for use with Intel(R) Parallel Studio XE toolchain and the Intel(R) MPI Library. |
| ohpc-intel-impi-perf-tools | Collection of performance tool builds for use with Intel(R) Parallel Studio XE compiler toolchain and the Intel(R) MPI runtime. |
| ohpc-intel-io-libs | Collection of IO library builds for use with Intel(R) Parallel Studio XE software suite. |
| ohpc-intel-mpich-io-libs | Collection of IO library builds for use with Intel(R) Parallel Studio XE software suite and MPICH runtime. |
| ohpc-intel-mpich-parallel-libs | Collection of parallel library builds for use with Intel(R) Parallel Studio XE toolchain and the MPICH runtime. |
| ohpc-intel-mpich-perf-tools | Collection of performance tool builds for use with Intel(R) Parallel Studio XE compiler toolchain and the MPICH runtime. |

Table 2 (cont): **Available OpenHPC Meta-packages**

| Group Name | Description |
|-----------------------------------|---|
| ohpc-intel-mvapich2-io-libs | Collection of IO library builds for use with Intel(R) Parallel Studio XE software suite and MVAPICH2 runtime. |
| ohpc-intel-mvapich2-parallel-libs | Collection of parallel library builds for use with Intel(R) Parallel Studio XE toolchain and the MVAPICH2 runtime. |
| ohpc-intel-mvapich2-perf-tools | Collection of performance tool builds for use with Intel(R) Parallel Studio XE compiler toolchain and the MVAPICH2 runtime. |
| ohpc-intel-openmpi4-io-libs | Collection of IO library builds for use with Intel(R) Parallel Studio XE software suite and OpenMPI runtime. |
| ohpc-intel-openmpi4-parallel-libs | Collection of parallel library builds for use with Intel(R) Parallel Studio XE toolchain and the OpenMPI runtime. |
| ohpc-intel-openmpi4-perf-tools | Collection of performance tool builds for use with Intel(R) Parallel Studio XE compiler toolchain and the OpenMPI runtime. |
| ohpc-intel-perf-tools | Collection of performance tool builds for use with Intel(R) Parallel Studio XE toolchain. |
| ohpc-intel-python3-libs | Collection of python3 related library builds for use with Intel(R) Parallel Studio XE toolchain. |
| ohpc-intel-runtimes | Collection of runtimes for use with Intel(R) Parallel Studio XE toolchain. |
| ohpc-intel-serial-libs | Collection of serial library builds for use with Intel(R) Parallel Studio XE toolchain. |
| ohpc-slurm-client | Collection of client packages for SLURM. |
| ohpc-slurm-server | Collection of server packages for SLURM. |
| ohpc-warewulf | Collection of base packages for Warewulf provisioning. |

What follows next in this Appendix is a series of tables that summarize the underlying RPM packages available in this OpenHPC release. These packages are organized by groupings based on their general functionality and each table provides information for the specific RPM name, version, brief summary, and the web URL where additional information can be obtained for the component. Note that many of the 3rd party community libraries that are pre-packaged with OpenHPC are built using multiple compiler and MPI families. In these cases, the RPM package name includes delimiters identifying the development environment for which each package build is targeted. Additional information on the OpenHPC package naming scheme is presented in §4.6. The relevant package groupings and associated Table references are as follows:

- Administrative tools (Table 3)
- Provisioning (Table 4)
- Resource management (Table 5)
- Compiler families (Table 6)
- MPI families (Table 7)
- Development tools (Table 8)
- Performance analysis tools (Table 9)
- Lustre (Table 10)
- IO Libraries (Table 11)
- Runtimes (Table 12)
- Serial/Threaded Libraries (Table 13)
- Parallel Libraries (Table 14)

Table 3: **Administrative Tools**

| RPM Package Name | Version | Info/URL |
|---|---------|---|
| conman-ohpc | 0.3.0 | ConMan: The Console Manager. http://dun.github.io/conman |
| docs-ohpc | 2.4.0 | OpenHPC documentation. https://github.com/openhpc/ohpc |
| examples-ohpc | 2.0 | Example source code and templates for use within OpenHPC environment. https://github.com/openhpc/ohpc |
| genders-ohpc | 1.27 | Static cluster configuration database. https://github.com/chaos/genders |
| lmod-defaults-gnu9-impi-ohpc lmod-defaults-gnu9-mpich-ofi-ohpc lmod-defaults-gnu9-mpich-ucx-ohpc lmod-defaults-gnu9-mvapich2-ohpc lmod-defaults-gnu9-openmpi4-ohpc lmod-defaults-intel-impi-ohpc lmod-defaults-intel-mpich-ohpc lmod-defaults-intel-mvapich2-ohpc lmod-defaults-intel-openmpi4-ohpc | 2.0 | OpenHPC default login environments. https://github.com/openhpc/ohpc |
| lmod-ohpc | 8.5.22 | Lua based Modules (lmod). https://github.com/TACC/Lmod |
| losf-ohpc | 0.56.0 | A Linux operating system framework for managing HPC clusters. https://github.com/hpcsi/losf |
| mrsh-ohpc | 2.12 | Remote shell program that uses munge authentication. https://github.com/chaos/mrsh |
| nhc-ohpc | 1.4.2 | LBNL Node Health Check. https://github.com/mej/nhc |
| ohpc-release | 2 | OpenHPC release files. https://github.com/openhpc/ohpc |
| pdsh-ohpc | 2.34 | Parallel remote shell program. https://github.com/chaos/pdsh |
| prun-ohpc | 2.2 | Convenience utility for parallel job launch. https://github.com/openhpc/ohpc |
| test-suite-ohpc | 2.4.0 | Integration test suite for OpenHPC. https://github.com/openhpc/ohpc/tests |

Table 4: **Provisioning**

| RPM Package Name | Version | Info/URL |
|---|---------|--|
| warewulf-common-ohpc | 3.9.0 | Scalable systems management suite for high performance clusters. http://warewulf.lbl.gov |
| warewulf-ipmi-ohpc-initramfs-aarch64 | 3.9.0 | Warewulf - Add IPMI to aarch64 initramfs. http://warewulf.lbl.gov |
| warewulf-ipmi-ohpc-initramfs-x86_64 | 3.9.0 | Warewulf - Add IPMI to x86_64 initramfs. http://warewulf.lbl.gov |
| warewulf-provision-ohpc-gpl-sources | 3.9.0 | Warewulf - GPL sources used in Warewulf provisioning. http://warewulf.lbl.gov |
| warewulf-cluster-ohpc | 3.9.0 | Warewulf - HPC cluster configuration. http://warewulf.lbl.gov |
| warewulf-ipmi-ohpc | 3.9.0 | Warewulf - IPMI support. http://warewulf.lbl.gov |
| warewulf-common-ohpc-localdb | 3.9.0 | Warewulf - Install local database server. http://warewulf.lbl.gov |
| warewulf-provision-ohpc | 3.9.0 | Warewulf - System provisioning core. http://warewulf.lbl.gov |
| warewulf-provision-ohpc-server | 3.9.0 | Warewulf - System provisioning server. http://warewulf.lbl.gov |
| warewulf-vnfs-ohpc | 3.9.0 | Warewulf - Virtual Node File System support. http://warewulf.lbl.gov |
| warewulf-provision-ohpc-server-ipxe-aarch64 | 3.9.0 | Warewulf - iPXE Bootloader for aarch64. http://warewulf.lbl.gov |
| warewulf-provision-ohpc-server-ipxe-x86_64 | 3.9.0 | Warewulf - iPXE Bootloader for x86_64. http://warewulf.lbl.gov |
| warewulf-provision-ohpc-initramfs-aarch64 | 3.9.0 | Warewulf - initramfs base for aarch64. http://warewulf.lbl.gov |
| warewulf-provision-ohpc-initramfs-x86_64 | 3.9.0 | Warewulf - initramfs base for x86_64. http://warewulf.lbl.gov |

Table 5: Resource Management

| RPM Package Name | Version | Info/URL |
|----------------------------|---------|---|
| magpie-ohpc | 2.5 | Scripts for running Big Data software in HPC environments. https://github.com/LLNL/magpie |
| openpbs-execution-ohpc | 20.0.1 | OpenPBS for an execution host. http://www.openpbs.org |
| openpbs-client-ohpc | 20.0.1 | OpenPBS for a client host. http://www.openpbs.org |
| openpbs-server-ohpc | 20.0.1 | OpenPBS for a server host. http://www.openpbs.org |
| pmix-ohpc | 3.1.4 | An extended/exascale implementation of PMI. https://pmix.github.io/pmix |
| slurm-devel-ohpc | 20.11.8 | Development package for Slurm. https://slurm.schedmd.com |
| slurm-example-configs-ohpc | 20.11.8 | Example config files for Slurm. https://slurm.schedmd.com |
| slurm-sview-ohpc | 20.11.8 | Graphical user interface to view and modify Slurm state. https://slurm.schedmd.com |
| slurm-pam_slurm-ohpc | 20.11.8 | PAM module for restricting access to compute nodes via Slurm. https://slurm.schedmd.com |
| slurm-perlapi-ohpc | 20.11.8 | Perl API to Slurm. https://slurm.schedmd.com |
| slurm-contribs-ohpc | 20.11.8 | Perl tool to print Slurm job state information. https://slurm.schedmd.com |
| slurm-ohpc | 20.11.8 | Slurm Workload Manager. https://slurm.schedmd.com |
| slurm-slurmd-ohpc | 20.11.8 | Slurm compute node daemon. https://slurm.schedmd.com |
| slurm-slurmetld-ohpc | 20.11.8 | Slurm controller daemon. https://slurm.schedmd.com |
| slurm-slurmdbd-ohpc | 20.11.8 | Slurm database daemon. https://slurm.schedmd.com |
| slurm-libpmi-ohpc | 20.11.8 | Slurm's implementation of the pmi libraries. https://slurm.schedmd.com |
| slurm-torque-ohpc | 20.11.8 | Torque/PBS wrappers for transition from Torque/PBS to Slurm. https://slurm.schedmd.com |
| slurm-openlava-ohpc | 20.11.8 | openlava/LSF wrappers for transition from OpenLava/LSF to Slurm. https://slurm.schedmd.com |

Table 6: Compiler Families

| RPM Package Name | Version | Info/URL |
|-----------------------------------|---------|--|
| gnu9-compilers-ohpc | 9.4.0 | The GNU C Compiler and Support Files. http://gcc.gnu.org |
| intel-oneapi-toolkit-release-ohpc | 2021 | Intel(R) oneAPI HPC Toolkit Repository Setup. https://github.com/openhpc/ohpc |
| intel-compilers-devel-ohpc | 2021 | OpenHPC compatability package for Intel(R) oneAPI HPC Toolkit. https://github.com/openhpc/ohpc |
| intel-psxe-compilers-devel-ohpc | 2021 | OpenHPC compatibility package for Intel(R) Parallel Studio XE. https://github.com/openhpc/ohpc |

Table 7: MPI Families / Communication Libraries

| RPM Package Name | Version | Info/URL |
|---------------------------|---------|--|
| intel-psxe-mpi-devel-ohpc | 2021 | OpenHPC compatibility package for Intel(R) MPI Library. https://github.com/openhpc/ohpc |
| intel-mpi-devel-ohpc | 2021 | OpenHPC compatibility package for Intel(R) oneAPI MPI Library. https://github.com/openhpc/ohpc |
| libfabric-ohpc | 1.13.0 | User-space RDMA Fabric Interfaces. http://www.github.com/ofiwg/libfabric |
| mpich-ofi-gnu9-ohpc | 3.4.2 | MPICH MPI implementation. |
| mpich-ofi-intel-ohpc | | http://www.mpich.org |
| mpich-ucx-gnu9-ohpc | 3.4.2 | MPICH MPI implementation. |
| mpich-ucx-intel-ohpc | | http://www.mpich.org |
| mvapich2-psm2-gnu9-ohpc | 2.3.4 | OSU MVAPICH2 MPI implementation. |
| mvapich2-psm2-intel-ohpc | | http://mvapich.cse.ohio-state.edu |
| mvapich2-gnu9-ohpc | 2.3.6 | OSU MVAPICH2 MPI implementation. |
| mvapich2-intel-ohpc | | http://mvapich.cse.ohio-state.edu |
| openmpi4-gnu9-ohpc | 4.1.1 | A powerful implementation of MPI/SHMEM. |
| openmpi4-intel-ohpc | | http://www.open-mpi.org |
| ucx-ohpc | 1.11.2 | UCX is a communication library implementing high-performance messaging. http://www.openucx.org |

Table 8: Development Tools

| RPM Package Name | Version | Info/URL |
|--|---------|--|
| EasyBuild-ohpc | 4.5.0 | Software build and installation framework. https://easybuilders.github.io/easybuild |
| automake-ohpc | 1.16.1 | A GNU tool for automatically creating Makefiles. http://www.gnu.org/software/automake |
| autoconf-ohpc | 2.69 | A GNU tool for automatically configuring source code. http://www.gnu.org/software/autoconf |
| cmake-ohpc | 3.21.3 | CMake is an open-source, cross-platform family of tools designed to build, test and package software. https://cmake.org |
| hwloc-ohpc | 2.5.0 | Portable Hardware Locality. http://www.open-mpi.org/projects/hwloc |
| libtool-ohpc | 2.4.6 | The GNU Portable Library Tool. http://www.gnu.org/software/libtool |
| python3-scipy-gnu9-mpich-ohpc python3-scipy-gnu9-mvapich2-ohpc python3-scipy-gnu9-openmpi4-ohpc | 1.5.1 | Scientific Tools for Python. http://www.scipy.org |
| python3-numpy-gnu9-ohpc python3-numpy-intel-ohpc | 1.19.5 | NumPy array processing for numbers, strings, records and objects. https://github.com/numpy/numpy |
| python3-mpi4py-gnu9-impi-ohpc python3-mpi4py-gnu9-mpich-ohpc python3-mpi4py-gnu9-mvapich2-ohpc python3-mpi4py-gnu9-openmpi4-ohpc python3-mpi4py-intel-impi-ohpc python3-mpi4py-intel-mpich-ohpc python3-mpi4py-intel-mvapich2-ohpc python3-mpi4py-intel-openmpi4-ohpc | 3.0.3 | Python bindings for the Message Passing Interface (MPI) standard. https://bitbucket.org/mpi4py/mpi4py |
| spack-ohpc | 0.17.0 | HPC software package management. https://github.com/LLNL/spack |
| valgrind-ohpc | 3.18.1 | Valgrind Memory Debugger. http://www.valgrind.org |

Table 9: Performance Analysis Tools

| RPM Package Name | Version | Info/URL |
|--|---------|--|
| dimemas-gnu9-impi-ohpc dimemas-gnu9-mpich-ohpc dimemas-gnu9-mvapich2-ohpc dimemas-gnu9-openmpi4-ohpc dimemas-intel-impi-ohpc dimemas-intel-mpich-ohpc dimemas-intel-mvapich2-ohpc dimemas-intel-openmpi4-ohpc | 5.4.2 | Dimemas tool. https://tools.bsc.es |
| extrae-gnu9-impi-ohpc extrae-gnu9-mpich-ohpc extrae-gnu9-mvapich2-ohpc extrae-gnu9-openmpi4-ohpc extrae-intel-impi-ohpc extrae-intel-mpich-ohpc extrae-intel-mvapich2-ohpc extrae-intel-openmpi4-ohpc | 3.7.0 | Extrae tool. https://tools.bsc.es |
| geopm-gnu9-impi-ohpc geopm-gnu9-mpich-ohpc geopm-gnu9-mvapich2-ohpc geopm-gnu9-openmpi4-ohpc geopm-intel-impi-ohpc geopm-intel-mpich-ohpc geopm-intel-mvapich2-ohpc geopm-intel-openmpi4-ohpc | 1.1.0 | Global Extensible Open Power Manager. https://geopm.github.io |
| imb-gnu9-impi-ohpc imb-gnu9-mpich-ohpc imb-gnu9-mvapich2-ohpc imb-gnu9-openmpi4-ohpc imb-intel-impi-ohpc imb-intel-mpich-ohpc imb-intel-mvapich2-ohpc imb-intel-openmpi4-ohpc | 2019.6 | Intel MPI Benchmarks (IMB). https://software.intel.com/en-us/articles/intel-mpi-benchmarks |
| likwid-gnu9-ohpc likwid-intel-ohpc | 5.0.1 | Performance tools for the Linux console. https://github.com/RRZE-HPC/likwid |
| msr-safe-ohpc | 1.6.0 | Allows safer access to model specific registers (MSRs). https://github.com/LLNL/msr-safe |
| omb-gnu9-impi-ohpc omb-gnu9-mpich-ohpc omb-gnu9-mvapich2-ohpc omb-gnu9-openmpi4-ohpc omb-intel-impi-ohpc omb-intel-mpich-ohpc omb-intel-mvapich2-ohpc omb-intel-openmpi4-ohpc | 5.8 | OSU Micro-benchmarks. http://mvapich.cse.ohio-state.edu/benchmarks |
| paraver-ohpc | 4.8.2 | Paraver. https://tools.bsc.es |
| papi-ohpc | 5.7.0 | Performance Application Programming Interface. http://icl.cs.utk.edu/papi |

Table 9 (cont): **Performance Analysis Tools**

| RPM Package Name | Version | Info/URL |
|--|---------|--|
| pdtoolkit-gnu9-ohpc pdtoolkit-intel-ohpc | 3.25.1 | PDT is a framework for analyzing source code. http://www.cs.uoregon.edu/Research/pdt |
| scalasca-gnu9-impi-ohpc scalasca-gnu9-mpich-ohpc scalasca-gnu9-mvapich2-ohpc scalasca-gnu9-openmpi4-ohpc scalasca-intel-impi-ohpc scalasca-intel-mpich-ohpc scalasca-intel-mvapich2-ohpc scalasca-intel-openmpi4-ohpc | 2.5 | Toolset for performance analysis of large-scale parallel applications. http://www.scalasca.org |
| scorep-gnu9-impi-ohpc scorep-gnu9-mpich-ohpc scorep-gnu9-mvapich2-ohpc scorep-gnu9-openmpi4-ohpc scorep-intel-impi-ohpc scorep-intel-mpich-ohpc scorep-intel-mvapich2-ohpc scorep-intel-openmpi4-ohpc | 6.0 | Scalable Performance Measurement Infrastructure for Parallel Codes. http://www.vi-hps.org/projects/score-p |
| tau-gnu9-impi-ohpc tau-gnu9-mpich-ohpc tau-gnu9-mvapich2-ohpc tau-gnu9-openmpi4-ohpc tau-intel-impi-ohpc tau-intel-mpich-ohpc tau-intel-mvapich2-ohpc tau-intel-openmpi4-ohpc | 2.29 | Tuning and Analysis Utilities Profiling Package. http://www.cs.uoregon.edu/research/tau/home.php |

Table 10: **Lustre**

| RPM Package Name | Version | Info/URL |
|--------------------|---------|--|
| lustre-client-ohpc | 2.12.7 | Lustre File System. https://wiki.whamcloud.com |

Table 11: IO Libraries

| RPM Package Name | Version | Info/URL |
|--|---------|--|
| adios-gnu9-impi-ohpc adios-gnu9-mpich-ohpc adios-gnu9-mvapich2-ohpc adios-gnu9-openmpi4-ohpc adios-intel-impi-ohpc adios-intel-mpich-ohpc adios-intel-mvapich2-ohpc adios-intel-openmpi4-ohpc | 1.13.1 | The Adaptable IO System (ADIOS). http://www.olcf.ornl.gov/center-projects/adios |
| hdf5-gnu9-ohpc hdf5-intel-ohpc | 1.10.8 | A general purpose library and file format for storing scientific data. http://www.hdfgroup.org/HDF5 |
| netcdf-cxx-gnu9-impi-ohpc netcdf-cxx-gnu9-mpich-ohpc netcdf-cxx-gnu9-mvapich2-ohpc netcdf-cxx-gnu9-openmpi4-ohpc netcdf-cxx-intel-impi-ohpc netcdf-cxx-intel-mpich-ohpc netcdf-cxx-intel-mvapich2-ohpc netcdf-cxx-intel-openmpi4-ohpc | 4.3.1 | C++ Libraries for the Unidata network Common Data Form. http://www.unidata.ucar.edu/software/netcdf |
| netcdf-fortran-gnu9-impi-ohpc netcdf-fortran-gnu9-mpich-ohpc netcdf-fortran-gnu9-mvapich2-ohpc netcdf-fortran-gnu9-openmpi4-ohpc netcdf-fortran-intel-impi-ohpc netcdf-fortran-intel-mpich-ohpc netcdf-fortran-intel-mvapich2-ohpc netcdf-fortran-intel-openmpi4-ohpc | 4.5.3 | Fortran Libraries for the Unidata network Common Data Form. http://www.unidata.ucar.edu/software/netcdf |
| netcdf-gnu9-impi-ohpc netcdf-gnu9-mpich-ohpc netcdf-gnu9-mvapich2-ohpc netcdf-gnu9-openmpi4-ohpc netcdf-intel-impi-ohpc netcdf-intel-mpich-ohpc netcdf-intel-mvapich2-ohpc netcdf-intel-openmpi4-ohpc | 4.7.4 | C Libraries for the Unidata network Common Data Form. http://www.unidata.ucar.edu/software/netcdf |
| phdf5-gnu9-impi-ohpc phdf5-gnu9-mpich-ohpc phdf5-gnu9-mvapich2-ohpc phdf5-gnu9-openmpi4-ohpc phdf5-intel-impi-ohpc phdf5-intel-mpich-ohpc phdf5-intel-mvapich2-ohpc phdf5-intel-openmpi4-ohpc | 1.10.8 | A general purpose library and file format for storing scientific data. http://www.hdfgroup.org/HDF5 |

Table 11 (cont): IO Libraries

| RPM Package Name | Version | Info/URL |
|--|---------|--|
| pnetcdf-gnu9-impi-ohpc pnetcdf-gnu9-mpich-ohpc pnetcdf-gnu9-mvapich2-ohpc pnetcdf-gnu9-openmpi4-ohpc pnetcdf-intel-impi-ohpc pnetcdf-intel-mpich-ohpc pnetcdf-intel-mvapich2-ohpc pnetcdf-intel-openmpi4-ohpc | 1.12.2 | A Parallel NetCDF library (PnetCDF). http://cucis.ece.northwestern.edu/projects/PnetCDF |
| sionlib-gnu9-impi-ohpc sionlib-gnu9-mpich-ohpc sionlib-gnu9-mvapich2-ohpc sionlib-gnu9-openmpi4-ohpc sionlib-intel-impi-ohpc sionlib-intel-mpich-ohpc sionlib-intel-mvapich2-ohpc sionlib-intel-openmpi4-ohpc | 1.7.4 | Scalable I/O Library for Parallel Access to Task-Local Files. http://www.fz-juelich.de/ias/jsc/EN/Expertise/Support/Software/SIONlib/_node.html |

Table 12: **Runtimes**

| RPM Package Name | Version | Info/URL |
|-------------------|---------|--|
| charliecloud-ohpc | 0.15 | Lightweight user-defined software stacks for high-performance computing. https://hpc.github.io/charliecloud |
| singularity-ohpc | 3.7.1 | Application and environment virtualization. https://www.sylabs.io/singularity |

Table 13: **Serial/Threaded Libraries**

| RPM Package Name | Version | Info/URL |
|---|---------|--|
| R-gnu9-ohpc | 4.1.2 | R is a language and environment for statistical computing and graphics (S-Plus like). http://www.r-project.org |
| gsl-gnu9-ohpc gsl-intel-ohpc | 2.7 | GNU Scientific Library (GSL). http://www.gnu.org/software/gsl |
| metis-gnu9-ohpc metis-intel-ohpc | 5.1.0 | Serial Graph Partitioning and Fill-reducing Matrix Ordering. http://glaros.dtc.umn.edu/gkhome/metis/metis/overview |
| openblas-gnu9-ohpc | 0.3.7 | An optimized BLAS library based on GotoBLAS2. http://www.openblas.net |
| plasma-gnu9-ohpc plasma-intel-ohpc | 2.8.0 | Parallel Linear Algebra Software for Multicore Architectures. https://bitbucket.org/icl/plasma |
| scotch-gnu9-ohpc scotch-intel-ohpc | 6.0.6 | Graph, mesh and hypergraph partitioning library. http://www.labri.fr/perso/pelegrin/scotch |
| superlu-gnu9-ohpc superlu-intel-ohpc | 5.2.1 | A general purpose library for the direct solution of linear equations. http://crd.lbl.gov/~xiaoye/SuperLU |

Table 14: **Parallel Libraries**

| RPM Package Name | Version | Info/URL |
|--|---------|---|
| boost-intel-impi-ohpc boost-intel-mpich-ohpc boost-intel-mvapich2-ohpc boost-intel-openmpi4-ohpc | 1.75.0 | Boost free peer-reviewed portable C++ source libraries. http://www.boost.org |
| boost-gnu9-impi-ohpc boost-gnu9-mpich-ohpc boost-gnu9-mvapich2-ohpc boost-gnu9-openmpi4-ohpc | 1.76.0 | |
| fftw-gnu9-impi-ohpc fftw-gnu9-mpich-ohpc fftw-gnu9-mvapich2-ohpc fftw-gnu9-openmpi4-ohpc | 3.3.8 | |
| hypre-gnu9-impi-ohpc hypre-gnu9-mpich-ohpc hypre-gnu9-mvapich2-ohpc hypre-gnu9-openmpi4-ohpc hypre-intel-impi-ohpc hypre-intel-mpich-ohpc hypre-intel-mvapich2-ohpc hypre-intel-openmpi4-ohpc | 2.18.1 | |
| mfem-gnu9-impi-ohpc mfem-gnu9-mpich-ohpc mfem-gnu9-mvapich2-ohpc mfem-gnu9-openmpi4-ohpc mfem-intel-impi-ohpc mfem-intel-mpich-ohpc mfem-intel-mvapich2-ohpc mfem-intel-openmpi4-ohpc | 4.3 | Lightweight, general, scalable C++ library for finite element methods. http://mfem.org |
| mumps-gnu9-impi-ohpc mumps-gnu9-mpich-ohpc mumps-gnu9-mvapich2-ohpc mumps-gnu9-openmpi4-ohpc mumps-intel-impi-ohpc mumps-intel-mpich-ohpc mumps-intel-mvapich2-ohpc mumps-intel-openmpi4-ohpc | 5.2.1 | A MULTifrontal Massively Parallel Sparse direct Solver. http://mumps.enseiht.fr |
| opencoarrays-gnu9-impi-ohpc opencoarrays-gnu9-mpich-ohpc opencoarrays-gnu9-mvapich2-ohpc opencoarrays-gnu9-openmpi4-ohpc | 2.9.2 | ABI to leverage the parallel programming features of the Fortran 2018 DIS. http://www.opencoarrays.org |
| petsc-gnu9-impi-ohpc petsc-gnu9-mpich-ohpc petsc-gnu9-mvapich2-ohpc petsc-gnu9-openmpi4-ohpc petsc-intel-impi-ohpc petsc-intel-mpich-ohpc petsc-intel-mvapich2-ohpc petsc-intel-openmpi4-ohpc | 3.16.1 | Portable Extensible Toolkit for Scientific Computation. http://www.mcs.anl.gov/petsc |

Table 14 (cont): **Parallel Libraries**

| RPM Package Name | Version | Info/URL |
|--|---------|---|
| ptscotch-gnu9-impi-ohpc ptscotch-gnu9-mpich-ohpc ptscotch-gnu9-mvapich2-ohpc ptscotch-gnu9-openmpi4-ohpc ptscotch-intel-impi-ohpc ptscotch-intel-mpich-ohpc ptscotch-intel-mvapich2-ohpc ptscotch-intel-openmpi4-ohpc | 6.0.6 | Graph, mesh and hypergraph partitioning library using MPI. http://www.labri.fr/perso/pelegrin/scotch |
| scalapack-gnu9-impi-ohpc scalapack-gnu9-mpich-ohpc scalapack-gnu9-mvapich2-ohpc scalapack-gnu9-openmpi4-ohpc scalapack-intel-impi-ohpc scalapack-intel-mpich-ohpc scalapack-intel-mvapich2-ohpc scalapack-intel-openmpi4-ohpc | 2.1.0 | A subset of LAPACK routines redesigned for heterogenous computing. http://www.netlib.org/lapack-dev |
| slepc-gnu9-impi-ohpc slepc-gnu9-mpich-ohpc slepc-gnu9-mvapich2-ohpc slepc-gnu9-openmpi4-ohpc slepc-intel-impi-ohpc slepc-intel-mpich-ohpc slepc-intel-mvapich2-ohpc slepc-intel-openmpi4-ohpc | 3.16.0 | A library for solving large scale sparse eigenvalue problems. http://slepc.upv.es |
| superlu_dist-gnu9-impi-ohpc superlu_dist-gnu9-mpich-ohpc superlu_dist-gnu9-mvapich2-ohpc superlu_dist-gnu9-openmpi4-ohpc superlu_dist-intel-impi-ohpc superlu_dist-intel-mpich-ohpc superlu_dist-intel-mvapich2-ohpc superlu_dist-intel-openmpi4-ohpc | 6.4.0 | A general purpose library for the direct solution of linear equations. http://crd-legacy.lbl.gov/~xiaoye/SuperLU |
| trilinos-gnu9-openmpi4-ohpc | 13.2.0 | A collection of libraries of numerical algorithms. https://trilinos.org |
| trilinos-gnu9-impi-ohpc trilinos-gnu9-mpich-ohpc trilinos-gnu9-mvapich2-ohpc trilinos-intel-impi-ohpc trilinos-intel-mpich-ohpc trilinos-intel-mvapich2-ohpc trilinos-intel-openmpi4-ohpc | 13.2.0 | A collection of libraries of numerical algorithms. https://trilinos.org |

F Package Signatures

All of the RPMs provided via the OpenHPC repository are signed with a GPG signature. By default, the underlying package managers will verify these signatures during installation to ensure that packages have not been altered. The RPMs can also be manually verified and the public signing key fingerprint for the latest repository is shown below:

Fingerprint: 5392 744D 3C54 3ED5 7847 65E6 8A30 6019 DA565C6C

The following command can be used to verify an RPM once it has been downloaded locally by confirming if the package is signed, and if so, indicating which key was used to sign it. The example below highlights usage for a local copy of the `docs-ohpc` package and illustrates how the *key ID* matches the fingerprint shown above.

```
[sms]# rpm --checksig -v docs-ohpc-*.rpm
docs-ohpc-2.0.0-72.1.ohpc.2.0.x86_64.rpm:
  Header V3 RSA/SHA1 Signature, key ID da565c6c: OK
  Header SHA256 digest: OK
  Header SHA1 digest: OK
  Payload SHA256 digest: OK
  V3 RSA/SHA1 Signature, key ID da565c6c: OK
  MD5 digest: OK
```