



OpenHPC (v2.2) Cluster Building Recipes

CentOS8.3 Base OS
xCAT/SLURM Edition for Linux* (x86_64)

Stateful Provisioning



Document Last Update: 2021-06-01
Document Revision: f1948ab41

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 xCAT and Provision Nodes with BOS | 8 |
| 3.1 | Enable xCAT repository for local use | 8 |
| 3.2 | Add provisioning services on <i>master</i> node | 8 |
| 3.3 | Complete basic xCAT setup for <i>master</i> node | 8 |
| 3.4 | Define <i>compute</i> image for provisioning | 9 |
| 3.4.1 | Build initial BOS image | 9 |
| 3.5 | Add compute nodes into xCAT database | 10 |
| 3.6 | Boot compute nodes | 11 |
| 4 | Install OpenHPC Components | 11 |
| 4.1 | Enable OpenHPC repository for local use | 11 |
| 4.2 | Installation template | 12 |
| 4.3 | Setup time synchronization service on <i>master</i> node | 12 |
| 4.4 | Add resource management services on <i>master</i> node | 12 |
| 4.5 | Optionally add InfiniBand support services on <i>master</i> node | 13 |
| 4.6 | Optionally add Omni-Path support services on <i>master</i> node | 14 |
| 4.6.1 | Add OpenHPC components to <i>compute</i> nodes | 15 |
| 4.6.2 | Customize system configuration | 16 |
| 4.6.3 | Additional Customization (<i>optional</i>) | 16 |
| 4.6.3.1 | Increase locked memory limits | 17 |
| 4.6.3.2 | Enable ssh control via resource manager | 17 |
| 4.6.3.3 | Add Lustre client | 17 |
| 4.6.3.4 | Add Nagios monitoring | 18 |
| 4.6.3.5 | Add ClusterShell | 19 |
| 4.6.3.6 | Add <i>genders</i> | 19 |
| 4.6.3.7 | Add Magpie | 19 |
| 4.6.3.8 | Add ConMan | 20 |
| 4.6.3.9 | Add NHC | 20 |
| 5 | Install OpenHPC Development Components | 20 |
| 5.1 | Development Tools | 21 |
| 5.2 | Compilers | 21 |
| 5.3 | MPI Stacks | 21 |
| 5.4 | Performance Tools | 22 |
| 5.5 | Setup default development environment | 22 |
| 5.6 | 3rd Party Libraries and Tools | 23 |
| 5.7 | Optional Development Tool Builds | 24 |
| 6 | Resource Manager Startup | 24 |

| | | |
|-------------------|--|-----------|
| 7 | Run a Test Job | 25 |
| 7.1 | Interactive execution | 26 |
| 7.2 | Batch execution | 27 |
| Appendices | | 28 |
| A | Installation Template | 28 |
| B | Upgrading OpenHPC Packages | 29 |
| C | Integration Test Suite | 30 |
| D | Customization | 32 |
| | D.1 Adding local Lmod modules to OpenHPC hierarchy | 32 |
| | D.2 Rebuilding Packages from Source | 33 |
| E | Package Manifest | 34 |
| F | Package Signatures | 49 |

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 stateful 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 CentOS8.3 and is subsequently configured to provision the remaining *compute* nodes with xCAT in a stateful 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
- `${domain_name}` # Local network domain name
- `${sms_eth_internal}` # Internal Ethernet interface on SMS
- `${internal_netmask}` # Subnet netmask for internal network
- `${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*")
- `${compute_prefix}` # Prefix for compute node names (e.g. "c")
- `${iso_path}` # Directory location of OS iso for xCAT install
- `${ohpc_repo_dir}` # Directory location of local OpenHPC repository mirror
- `${epel_repo_dir}` # Directory location of local EPEL repository mirror

Optional:

- `${sysmgmtd_host}` # BeeGFS System Management host name
- `${mgs_fs_name}` # Lustre MGS mount name
- `${sms_ipoib}` # IPoIB address for SMS server
- `${ipoib_netmask}` # Subnet netmask for internal IPoIB
- `${c_ipoib[0]}, ${c_ipoib[1]}, ...` # IPoIB addresses for computes
- `${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 CentOS8.3 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 CentOS8.3 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 xCAT, 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 xCAT and Provision Nodes with BOS

Stateful installation is accomplished in two steps: First, a generic OS image is installed on *compute* nodes and then, once the nodes are up and running, OpenHPC components are added to both the SMS and the nodes at the same time.

3.1 Enable xCAT repository for local use

To begin, enable use of the public xCAT repository by adding it to the local list of available package repositories. This also requires network access from your *master* server to the internet, or alternatively, that the repository be mirrored locally. In this case, we use the network.

```
[sms]# yum -y install yum-utils
[sms]# wget -P /etc/yum.repos.d https://xcat.org/files/xcat/repos/yum/latest/xcat-core/xcat-core.repo
```

xCAT has a number of dependencies that are required for installation that are housed in separate public repositories for various distributions. To enable for local use, issue the following:

```
[sms]# yum -y install centos-release-stream
[sms]# wget -P /etc/yum.repos.d https://xcat.org/files/xcat/repos/yum/xcat-dep/rh8/x86_64/xcat-dep.repo
```

3.2 Add provisioning services on *master* node

With xCAT repository enabled, issue the following install the provisioning service on *master* node

```
[sms]# yum -y install xCAT
# enable xCAT tools for use in current shell
[sms]# . /etc/profile.d/xcat.sh
```

3.3 Complete basic xCAT setup for *master* node

At this point, all of the packages necessary to use xCAT on the *master* host should be installed. Next, we enable support for local provisioning using a second private interface (refer to Figure 1) and register this network interface with xCAT.

```
# 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}

# Register internal provisioning interface with xCAT for DHCP
[sms]# chdef -t site dhcpinterfaces="xcatmn|${sms_eth_internal}"
```


xCAT requires a network domain name specification for system-wide name resolution. This value can be set to match your local DNS schema or given a unique identifier such as “local”. In this recipe, we leverage the `$domain_name` variable to define as follows:

```
[sms]# chdef -t site domain=${domain_name}
```

3.4 Define *compute* image for provisioning

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

3.4.1 Build initial BOS image

The following steps illustrate the process to build a minimal, default image for use with xCAT. To begin, you will first need to have a local copy of the ISO image available for the underlying OS. In this recipe, the relevant ISO image is `CentOS-8.3.2011-x86_64-dvd1.iso` (available from the CentOS mirrors). We initialize the image creation process using the `copycds` command assuming that the necessary ISO image is available locally in `${iso_path}` as follows:

```
[sms]# copycds ${iso_path}/CentOS-8.3.2011-x86_64-dvd1.iso
```

Once completed, several OS images should be available for use within xCAT. These can be queried via:

```
# Query available images
[sms]# lsdef -t osimage
centos8-x86_64-install-compute  (osimage)
centos8-x86_64-netboot-compute  (osimage)
centos8-x86_64-stateful-mgmtnode (osimage)
centos8-x86_64-statelite-compute (osimage)
```

In this example, we leverage the stateful (*install*) image for provisioning of *compute* nodes. The other images, *netboot-compute* and *statelite-compute*, are for stateless and statelite installations, respectively. OpenHPC provides a stateless xCAT recipe, follow that guide if you are interested in a stateless install. Statelite installation is an intermediate type of install, in which a limited number of files and directories persist across reboots. Please consult available xCAT [documentation](#) if interested in this type of install.

3.5 Add compute nodes into xCAT database

Next, we add *compute* nodes and define their properties as objects in xCAT database. These hosts are grouped logically into a group named *compute* to facilitate group-level commands used later in the recipe. Note the use of variable names for the desired compute hostnames, node IPs, MAC addresses, and BMC login credentials, which should be modified to accommodate local settings and hardware. To enable serial console access via xCAT, serialport and serialspeed properties are also defined.

```
# Define nodes as objects in xCAT database
[sms]# for ((i=0; i<$num_computes; i++)) ; do
    mkdef -t node ${c_name[$i]} groups=compute,all ip=${c_ip[$i]} mac=${c_mac[$i]} netboot=xnba \
    arch=x86_64 bmc=${c_bmc[$i]} bmcusername=${bmc_username} bmcpassword=${bmc_password} \
    mgt=ipmi serialport=0 serialspeed=115200
done
```

Tip

Defining nodes one-by-one, as done above, is only efficient for a small number of nodes. For larger node counts, xCAT provides capabilities for automated detection and configuration. Consult the [xCAT Hardware Discovery & Define Node Guide](#). Alternatively, confluent, a tool related to xCAT, also has robust discovery capabilities and can be used to [detect and auto-configure](#) compute hosts.

Next, we set root password for *compute* nodes. Note that xCAT sets up *ssh* keys to enable password-less logins for root across the cluster, but the password is required for kickstarting the stateful install.

```
[sms]# chtab key=system passwd.username=root passwd.password=`openssl rand -base64 12`
```

If enabling *optional* IPoIB functionality (e.g. to support Lustre over InfiniBand), additional settings are required to define the IPoIB network with xCAT and specify desired IP settings for each compute. This can be accomplished as follows for the *ib0* interface:

```
# Define ib0 netmask
[sms]# chdef -t network -o ib0 mask=$ipoib_netmask net=${c_ipoib[0]}

# Enable secondary NIC configuration
[sms]# chdef compute -p postbootscripts=confignics

# Register desired IPoIB IPs per compute
[sms]# for ((i=0; i<$num_computes; i++)) ; do
    chdef ${c_name[i]} nicips.ib0=${c_ipoib[i]} nictypes.ib0="InfiniBand" nicnetworks.ib0=ib0
done
```

With the desired compute nodes and domain identified, the remaining steps in the provisioning configuration process are to define the provisioning mode and image for the *compute* group and use xCAT commands to complete configuration for network services like DNS and DHCP. These tasks are accomplished as follows:

```
# Complete network service configurations
[sms]# makehosts
[sms]# makenetworks
[sms]# makedhcp -n
[sms]# makedns -n

# Associate desired provisioning image for computes
[sms]# nodeset compute osimage=centos8-x86_64-install-compute
```

3.6 Boot compute nodes

Prior to booting the *compute* hosts, we configure them to use PXE as their next boot mode. After the initial PXE, ensuing boots will return to using the default boot device specified in the BIOS.

```
[sms]# rsetboot compute net
```

At this point, the *master* server should be able to boot the newly defined compute nodes. This is done by using the **rpower** xCAT command leveraging IPMI protocol set up during the the *compute* node definition in § 3.5. The following power cycles each of the desired hosts.

```
[sms]# rpower compute reset
```

Once kicked off, the boot process should take about 5 to 10 minutes. You can monitor the provisioning by using the **rcons** command, which displays serial console for a selected node. Note that the escape sequence is CTRL-e c . typed sequentially.

Successful provisioning can be verified by a parallel command on the compute nodes. The default install provides two such tools: xCAT-provided **psh** command, which uses xCAT node names and groups, and **pdsh**, which works independently. For example, to run a command on the newly imaged compute hosts using **psh**, execute the following:

```
[sms]# psh compute uptime
c1: 12:56:50 up 14 min, 0 users, load average: 0.00, 0.01, 0.04
c2: 12:56:50 up 13 min, 0 users, load average: 0.00, 0.02, 0.05
c3: 12:56:50 up 14 min, 0 users, load average: 0.00, 0.02, 0.05
c4: 12:56:50 up 14 min, 0 users, load average: 0.00, 0.01, 0.04
```

Note that the equivalent **pdsh** command is **pdsh -w c[1-4] uptime**.

4 Install OpenHPC Components

With the BOS installed and booted on both the *master* server and the *compute* nodes, the next step is to add desired OpenHPC packages into the cluster. The following subsections highlight this process.

4.1 Enable OpenHPC repository for local use

OpenHPC provides an online package repository, but here we setup and use a local mirror on the SMS host for OpenHPC package installation. This is preferable for a stateful install, because *compute* nodes can access it on the local network.

Assuming that the *master* has outside network access, the mirror is created by downloading a tarball of the repository, unpacking it and running a provided script to configure the package manager.

```
# Download OHPC tarball for local installation
[sms]# wget http://repos.openhpc.community/dist/2.2/OpenHPC-2.2.CentOS_8.x86_64.tar
# Create directory if necessary
[sms]# mkdir -p $ohpc_repo_dir
#Unpack
[sms]# tar xvf OpenHPC-2.2.CentOS_8.x86_64.tar -C $ohpc_repo_dir
#Create local mirror repository
[sms]# $ohpc_repo_dir/make_repo.sh
```

In addition to the OpenHPC and xCAT package repositories, the *master* host also requires access to the standard base OS distro repositories in order to resolve necessary dependencies. For CentOS8.3, the

requirements are to have access to both the base OS and EPEL repositories for which mirrors are freely available online:

- CentOS-8 - Base 8.3.2011 (e.g. http://mirror.centos.org/centos-8/8/BaseOS/x86_64/os)
- EPEL 8 (e.g. http://download.fedoraproject.org/pub/epel/8/Everything/x86_64)

The public EPEL repository is enabled by installing `epel-release` package. Note that this requires the CentOS Extras repository, which is shipped with CentOS and is enabled by default.

```
[sms]# yum -y install epel-release
```

Now OpenHPC packages can be installed. To add the base package on the SMS issue the following

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

4.2 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.

4.3 Setup time synchronization service on *master* node

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
```

4.4 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.

4.5 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"

# Load IB drivers
[sms]# systemctl start rdma
```

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 CentOS8.3 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 §4.6.3.3 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
```

4.6 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

# Load RDMA services
[sms]# systemctl start rdma
```

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 CentOS8.3 provides the `opa-fm` package should you choose to run it on the *master* node.

4.6.1 Add OpenHPC components to *compute* nodes

The next step is adding OpenHPC components to the *compute* nodes that at this point are running basic OSes. This process will leverage two xCAT-provided commands: `psh` to run `yum` installer on all the nodes in parallel and `xdcp` to distribute configuration files from the SMS to the *compute* nodes.

To do this, repositories on the *compute* nodes need to be configured properly.

xCAT has automatically setup an OS repository on the SMS and configured the nodes to use it, but it has also enabled online OS repositories, which will cause `yum` failures if there is no connection to the public internet. To avoid this, we disable the online repositories:

```
# Install yum-utils to have yum-config-manager available
[sms]# psh compute yum --setopt=*.skip_if_unavailable=1 -y install yum-utils perl

# Enable CentOS-PowerTools repository
[sms]# psh compute yum-config-manager --enable PowerTools
```

Next, the nodes are configured to use the local OpenHPC repository setup in on the SMS in §4.1

```
# Add OpenHPC repo mirror hosted on SMS
[sms]# psh compute yum-config-manager --add-repo=http://$sms_ip/$ohpc_repo_dir/OpenHPC.local.repo
# Replace local path with SMS URL
[sms]# psh compute "perl -pi -e 's/file:\/\@PATH\/@http:\/\@\/$sms_ip\/\"${ohpc_repo_dir}\/\@\/\"/s' \
/etc/yum.repos.d/OpenHPC.local.repo"
```

The *compute* nodes also need access to the EPEL repository, a required dependency for OpenHPC packages. A straightforward way to accomplish this is to host a local mirror on the *master* host. However, the repository is very large (>10 GB) containing thousands of packages, while only 3 (`fping`, `libconfuse`, and `libunwind`) are required for packages installed this guide. Therefore, here we only download the required packages and put them into a repository accessible to the nodes. Users wanting to mirror the entire EPEL repository instead can take advantage of the `reposync` utility and should expect a download time of over an hour depending on their network connection and a choice of EPEL mirror.

```
# Create directory holding EPEL packages
[sms]# mkdir -p $epel_repo_dir
[sms]# yum -y install yum-utils createrepo
# Download required EPEL packages
[sms]# yumdownloader --destdir $epel_repo_dir fping libconfuse libunwind
# Create a repository from the downloaded packages
[sms]# createrepo $epel_repo_dir

# Setup access to the repo from the compute hosts
[sms]# psh compute yum-config-manager --add-repo=http://$sms_ip/$epel_repo_dir
[sms]# psh compute "echo gpgcheck=0 >> /etc/yum.repos.d/${sms_ip}-${epel_repo_dir}\/\@\/\"}.repo"
```

Additionally, a workaround is needed for OpenHPC documentation files, which are installed into a read-only NFS share `/opt/ohpc/pub`. Any package attempting to write to that directory will fail to install. The following prevents that by directing `rpm` not to install documentation files on the *compute* nodes:

```
[sms]# psh compute echo -e %_excludedocs 1 \>\> ~/.rpmmacros
```

Now OpenHPC and other cluster-related software components can be installed on the nodes. The first step is to install a base compute package:

```
# Install compute node base meta-package
[sms]# psh compute yum -y install ohpc-base-compute
```

Next, we can include additional components:

```
# Add Slurm client support meta-package
[sms]# psh compute yum -y install ohpc-slurm-client

# Add Network Time Protocol (NTP) support
[sms]# psh compute yum -y install ntp

# Add kernel drivers
[sms]# psh compute yum -y install kernel

# Include modules user environment
[sms]# psh compute yum -y install --enablerepo=powertools lmod-ohpc
```

```
# Optionally add IB support and enable
[sms]# psh compute yum -y groupinstall "InfiniBand Support"
[sms]# psh compute systemctl enable rdma
[sms]# psh compute systemctl start rdma
```

4.6.2 Customize system configuration

Here we set up 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.

```
# Disable /tftpboot and /install export entries
[sms]# perl -pi -e "s|/tftpboot|#/tftpboot|" /etc/exports
[sms]# perl -pi -e "s|/install|#/install|" /etc/exports

# 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
[sms]# exportfs -a
[sms]# systemctl restart nfs-server
[sms]# systemctl enable nfs-server

# Create NFS client mounts of /home and /opt/ohpc/pub on compute hosts
[sms]# psh compute echo \
    "\"${sms_ip}:/home /home nfs nfsvers=3,nodev,nosuid 0 0\" \">> /etc/fstab
[sms]# psh compute echo \
    "\"${sms_ip}:/opt/ohpc/pub /opt/ohpc/pub nfs nfsvers=3,nodev 0 0\" \">> /etc/fstab
[sms]# psh compute systemctl restart nfs

# Mount NFS shares
[sms]# psh compute mount /home
[sms]# psh compute mkdir -p /opt/ohpc/pub
[sms]# psh compute mount /opt/ohpc/pub
```

4.6.3 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
- 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.

4.6.3.1 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 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 compute nodes:

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

4.6.3.2 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 on *compute* nodes, issue the following:

```
[sms]# psh compute echo "\"account required pam_slurm.so\"" \>\> /etc/pam.d/sshd
```

4.6.3.3 Add Lustre client To add Lustre client support on the cluster, it 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
```

```
# Enable lustre on compute nodes
[sms]# psh compute yum -y install lustre-client-ohpc

# Create mount point and file system mount
[sms]# psh compute mkdir /mnt/lustre
[sms]# psh compute echo \
    "\"${mgs_fs_name} /mnt/lustre lustre defaults,_netdev,localflock,retry=2 0 0\"" \>\> /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* nodes.

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

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

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

# Mount on compute nodes
[sms]# psh compute mount /mnt/lustre
```

4.6.3.4 Add Nagios monitoring Nagios is an open source infrastructure monitoring package that monitors servers, switches, applications, and services and offers user-defined alerting facilities. As provided by OpenHPC, it consists of a base monitoring daemon and a set of plug-ins for monitoring various aspects of an HPC cluster. 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.

```
# Install Nagios meta-package on master host
[sms]# yum -y install ohpc-nagios

# Install plugins on compute nodes
[sms]# psh compute yum -y install nagios-plugins-all-ohpc nrpe-ohpc

# Enable and configure NRPE on compute nodes
[sms]# psh compute systemctl enable nrpe
[sms]# psh compute perl -pi -e "\"s/^allowed_hosts=/\# allowed_hosts=\" /etc/nagios/nrpe.cfg
[sms]# psh compute echo "\"nrpe 5666/tcp \# NRPE\"" \>\> /etc/services
[sms]# psh compute echo "\"nrpe : ${sms_ip} : ALLOW\"" \>\> /etc/hosts.allow
[sms]# psh compute echo "\"nrpe : ALL : DENY\"" \>\> /etc/hosts.allow
[sms]# psh compute "/usr/sbin/useradd -c \"NRPE user for the NRPE service\" -d /var/run/nrpe \
-r -g nrpe -s /sbin/nologin nrpe"
[sms]# psh compute /usr/sbin/groupadd -r nrpe

# Configure remote services to test on compute nodes
[sms]# mv /etc/nagios/conf.d/services.cfg.example /etc/nagios/conf.d/services.cfg

# Define compute nodes as hosts to monitor
[sms]# mv /etc/nagios/conf.d/hosts.cfg.example /etc/nagios/conf.d/hosts.cfg
[sms]# for ((i=0; i<$num_computes; i++)) ; do
    perl -pi -e "s/HOSTNAME${(i+1)}/${c_name[i]}/ || s/HOST${(i+1)}_IP/${c_ip[i]}/" \
    /etc/nagios/conf.d/hosts.cfg
done
```

```
# 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]# psh compute "echo command[check_ssh]=/usr/lib64/nagios/plugins/check_ssh localhost \
>> /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]# chkconfig nagios on
[sms]# systemctl start nagios
[sms]# chmod u+s `which ping`

# Start on compute hosts
[sms]# psh compute systemctl start nrpe
```

4.6.3.5 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
```

4.6.3.6 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
```

4.6.3.7 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
```

4.6.3.8 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<$num_compute; i++)) ; do
    echo -n 'CONSOLE name='${c_name[$i]}' dev="ipmi:'${c_bmc[$i]}'" '
    echo 'ipmiopts="'U:${bmc_username},P:${IPMI_PASSWORD:-undefined},W:solpayloadsize'"
done >> /etc/conman.conf

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

Note that additional options are typically necessary to enable serial console output. These are setup during the node registration process in §3.5

4.6.3.9 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]# psh compute yum -y 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
```

5 Install OpenHPC Development Components

The install procedure outlined in §4 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.

5.1 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
```

5.2 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
```

5.3 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 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 §4.5 and §4.6.3, 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
```

5.4 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
```

5.5 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

5.6 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 §5.3.

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 §4.6.2, 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
```

5.7 Optional Development Tool Builds

In addition to the 3rd party development libraries built using the open source toolchains mentioned in §5.6, OpenHPC also provides *optional* compatible builds for use with the compilers and MPI stack included in newer versions of the Intel® Parallel Studio XE software suite. 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, the Parallel Studio software suite must be obtained and installed separately. Once installed locally, the OpenHPC compatible packages can be installed using standard package manager semantics. Note that licenses are provided free of charge for many categories of use. In particular, licenses for compilers and developments tools are provided at no cost to academic researchers or developers contributing to open-source software projects. More information on this program can be found at:

<https://software.intel.com/en-us/qualify-for-free-software>

Tip

As noted in §4.6.2, the default installation path for OpenHPC (`/opt/ohpc/pub`) is exported over NFS from the *master* to the compute nodes, but the Parallel Studio installer defaults to a path of `/opt/intel`. To make the Intel® compilers available to the compute nodes one must either customize the Parallel Studio installation path to be within `/opt/ohpc/pub`, or alternatively, 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 Intel® Parallel Studio, issue the following:

```
# Install OpenHPC compatibility packages (requires prior installation of Parallel Studio)
[sms]# yum -y install intel-compilers-devel-ohpc
[sms]# yum -y install intel-mpi-devel-ohpc
```

```
# 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-io-libs
[sms]# yum -y install ohpc-intel-perf-tools
[sms]# yum -y install ohpc-intel-python3-libs
[sms]# yum -y install ohpc-intel-runtimes
[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
```

6 Resource Manager Startup

In section §4, the Slurm resource manager was installed and configured on the *master* host. Slurm clients were also installed, but have not been configured yet. To do so, Slurm and **munge** configuration files need to be copied to the nodes. This can be accomplished as follows:


```
[sms]# xdcp compute /etc/slurm/slurm.conf /etc/slurm/slurm.conf
[sms]# xdcp compute /etc/munge/munge.key /etc/munge/munge.key
```

With Slurm configured, 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 slurmd
[sms]# systemctl start munge
[sms]# systemctl start slurmd

# Start slurm clients on compute hosts
[sms]# psh compute systemctl enable munge
[sms]# psh compute systemctl enable slurmd
[sms]# psh compute systemctl start munge
[sms]# psh compute systemctl start slurmd
```

After this, check status of the nodes within Slurm by using the `sinfo` command. All compute nodes should be in an *idle* state (without asterisk). If the state is reported as *unknown*, the following might help:

```
[sms]# scontrol update partition=normal state=idle
```

In case of additional Slurm issues, ensure that the configuration file fits your hardware and that it is identical across the nodes. Also, verify that Slurm user id is the same on the SMS and *compute* nodes. You may also consult [Slurm Troubleshooting Guide](#).

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
```

Next, the user’s credentials need to be distributed across the cluster. xCAT’s `xdcp` has a merge functionality that adds new entries into credential files on *compute* nodes:

```
# Create a sync file for pushing user credentials to the nodes
[sms]# echo "MERGE:" > syncusers
[sms]# echo "/etc/passwd -> /etc/passwd" >> syncusers
[sms]# echo "/etc/group -> /etc/group" >> syncusers
[sms]# echo "/etc/shadow -> /etc/shadow" >> syncusers
# Use xCAT to distribute credentials to nodes
[sms]# xdcp compute -F syncusers
```

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 ~]$ srun -n 8 -N 2 --pty /bin/bash

[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 | srun -pty /bin/bash |
| 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/centos8/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/centos8/x86_64/xcat/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 §4.1 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.

The initial step, when using a local mirror as described in §4.1, is downloading a new tarball from <http://build.openhpc.community/dist>. After this, move (or remove) the old repository directory, unpack the new tarball and run the `make_repo.sh` script. In the following, substitute `x.x.x` with the desired version number.

```
[sms]# wget http://repos.openhpc.community/dist/x.x.x/OpenHPC-x.x.x.CentOS_8.x86_64.tar
[sms]# mv $ohpc_repo_dir "$ohpc_repo_dir"_old
[sms]# mkdir -p $ohpc_repo_dir
[sms]# tar xvf OpenHPC-x.x.x.CentOS_8.x86_64.tar -C $ohpc_repo_dir
[sms]# $ohpc_repo_dir/make_repo.sh
```

The *compute* hosts will pickup new packages automatically as long as the repository location on the SMS stays the same. If you prefer to use a different location, setup repositories on the nodes as described in §4.1.

The actual update is performed as follows:

1. (Optional) Ensure repo metadata is current (on head and compute nodes.) 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]# psh compute yum clean expire-cache
```

2. Upgrade master (SMS) node

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

3. Upgrade packages on compute nodes

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

Note that to update running services such as a resource manager a service restart (or a full reboot) is required.

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. 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
hybre/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 §5.6. The relevant package groupings and associated Table references are as follows:

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

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 | 2.2.0 | 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 | 8.2.10 | 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 | 2.1 | Convenience utility for parallel job launch. https://github.com/openhpc/ohpc |
| test-suite-ohpc | 2.1.0 | Integration test suite for OpenHPC. https://github.com/openhpc/ohpc/tests |

Table 4: **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 | 20.0.0 | OpenPBS for an execution host. http://www.openpbs.org |
| openpbs-client-ohpc | 20.0.0 | OpenPBS for a client host. http://www.openpbs.org |
| openpbs-server-ohpc | 20.0.0 | 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 | 20.11.7 | Development package for Slurm. https://slurm.schedmd.com |
| slurm-example-configs-ohpc | 20.11.7 | Example config files for Slurm. https://slurm.schedmd.com |
| slurm-sview-ohpc | 20.11.7 | Graphical user interface to view and modify Slurm state. https://slurm.schedmd.com |
| slurm-pam_slurm-ohpc | 20.11.7 | PAM module for restricting access to compute nodes via Slurm. https://slurm.schedmd.com |
| slurm-perlapi-ohpc | 20.11.7 | Perl API to Slurm. https://slurm.schedmd.com |
| slurm-contribs-ohpc | 20.11.7 | Perl tool to print Slurm job state information. https://slurm.schedmd.com |
| slurm-ohpc | 20.11.7 | Slurm Workload Manager. https://slurm.schedmd.com |
| slurm-slurmd-ohpc | 20.11.7 | Slurm compute node daemon. https://slurm.schedmd.com |
| slurm-slurmdctld-ohpc | 20.11.7 | Slurm controller daemon. https://slurm.schedmd.com |
| slurm-slurmdbd-ohpc | 20.11.7 | Slurm database daemon. https://slurm.schedmd.com |
| slurm-libpmi-ohpc | 20.11.7 | Slurm's implementation of the pmi libraries. https://slurm.schedmd.com |
| slurm-torque-ohpc | 20.11.7 | Torque/PBS wrappers for transition from Torque/PBS to Slurm. https://slurm.schedmd.com |
| slurm-openlava-ohpc | 20.11.7 | openlava/LSF wrappers for transition from OpenLava/LSF to Slurm. https://slurm.schedmd.com |

Table 5: **Compiler Families**

| RPM Package Name | Version | Info/URL |
|----------------------------|---------|---|
| gnu9-compilers-ohpc | 9.3.0 | The GNU C Compiler and Support Files. http://gcc.gnu.org |
| intel-compilers-devel-ohpc | 2020 | OpenHPC compatibility package for Intel(R) Parallel Studio XE. https://github.com/openhpc/ohpc |

Table 6: MPI Families / Communication Libraries

| RPM Package Name | Version | Info/URL |
|--|---------|--|
| intel-mpi-devel-ohpc | 2020 | OpenHPC compatibility package for Intel(R) MPI Library. https://github.com/openhpc/ohpc |
| libfabric-ohpc | 1.11.2 | User-space RDMA Fabric Interfaces. http://www.github.com/ofiwg/libfabric |
| mpich-ofi-gnu9-ohpc mpich-ofi-intel-ohpc | 3.3.2 | MPICH MPI implementation. http://www.mpich.org |
| mpich-ucx-gnu9-ohpc mpich-ucx-intel-ohpc | 3.3.2 | MPICH MPI implementation. http://www.mpich.org |
| mvapich2-gnu9-ohpc mvapich2-intel-ohpc mvapich2-psm2-gnu9-ohpc mvapich2-psm2-intel-ohpc | 2.3.4 | OSU MVAPICH2 MPI implementation. http://mvapich.cse.ohio-state.edu |
| openmpi4-gnu9-ohpc openmpi4-intel-ohpc | 4.0.5 | A powerful implementation of MPI/SHMEM. http://www.open-mpi.org |
| ucx-ohpc | 1.9.0 | UCX is a communication library implementing high-performance messaging. http://www.openucx.org |

Table 7: Development Tools

| RPM Package Name | Version | Info/URL |
|--|---------|--|
| EasyBuild-ohpc | 4.3.4 | 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 | 3.19.4 | CMake is an open-source, cross-platform family of tools designed to build, test and package software. https://cmake.org |
| hwloc-ohpc | 2.1.0 | 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 | 1.19.0 | 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 | 0.15.0 | HPC software package management. https://github.com/LLNL/spack |
| valgrind-ohpc | 3.16.1 | Valgrind Memory Debugger. http://www.valgrind.org |

Table 8: 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 | 1.4.0 | 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 | 5.6.2 | 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 8 (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 9: **Lustre**

| RPM Package Name | Version | Info/URL |
|--------------------|---------|--|
| lustre-client-ohpc | 2.12.6 | Lustre File System. https://wiki.whamcloud.com |

Table 10: 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 | 1.10.6 | 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 | 4.5.2 | 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 | 4.7.3 | 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 | 1.10.6 | A general purpose library and file format for storing scientific data. http://www.hdfgroup.org/HDF5 |

Table 10 (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 | 1.12.1 | 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 11: **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 12: **Serial/Threaded Libraries**

| RPM Package Name | Version | Info/URL |
|---|---------|--|
| R-gnu9-ohpc | 3.6.3 | 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 | 2.6 | 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 13: **Parallel Libraries**

| RPM Package Name | Version | Info/URL |
|--|---------|---|
| boost-gnu9-impi-ohpc boost-gnu9-mpich-ohpc boost-gnu9-mvapich2-ohpc boost-gnu9-openmpi4-ohpc 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 |
| fftw-gnu9-impi-ohpc fftw-gnu9-mpich-ohpc fftw-gnu9-mvapich2-ohpc fftw-gnu9-openmpi4-ohpc | 3.3.8 | A Fast Fourier Transform library. http://www.fftw.org |
| 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 | Scalable algorithms for solving linear systems of equations. http://www.llnl.gov/casc/hypre |
| 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 | 4.2 | 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 | 3.14.4 | Portable Extensible Toolkit for Scientific Computation. http://www.mcs.anl.gov/petsc |

Table 13 (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/pelegriin/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 | 3.14.2 | 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 | 6.1.1 | A general purpose library for the direct solution of linear equations. http://crd-legacy.lbl.gov/~xiaoye/SuperLU |
| trilinos-gnu9-impi-ohpc trilinos-gnu9-mpich-ohpc trilinos-gnu9-mvapich2-ohpc trilinos-gnu9-openmpi4-ohpc trilinos-intel-impi-ohpc trilinos-intel-mpich-ohpc trilinos-intel-mvapich2-ohpc trilinos-intel-openmpi4-ohpc | 13.0.0 | 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
```