

---

# **Omnia**

***Release 1.5***

**dell/omnia**

**Oct 18, 2023**



# CONTENTS

|          |   |            |
|----------|---|------------|
| <b>1</b> | <b>Omnia: Overview</b>                              | <b>3</b>   |
| 1.1      | Architecture . . . . .                              | 4          |
| 1.2      | New Features . . . . .                              | 6          |
| 1.3      | Releases . . . . .                                  | 6          |
| 1.4      | Support Matrix . . . . .                            | 9          |
| 1.5      | Network Topologies . . . . .                        | 15         |
| 1.6      | Blogs about Omnia . . . . .                         | 19         |
| 1.7      | What Omnia does . . . . .                           | 19         |
| <b>2</b> | <b>Quick Installation Guide</b>                     | <b>21</b>  |
| 2.1      | Running prereq.sh . . . . .                         | 21         |
| 2.2      | Installing the provision tool . . . . .             | 22         |
| 2.3      | Creating node inventory . . . . .                   | 57         |
| 2.4      | Configuring the cluster . . . . .                   | 57         |
| 2.5      | Adding new nodes . . . . .                          | 80         |
| 2.6      | Re-provisioning the cluster . . . . .               | 83         |
| 2.7      | Configuring switches . . . . .                      | 84         |
| 2.8      | Configuring storage . . . . .                       | 93         |
| 2.9      | Running HPC benchmarks on omnia clusters . . . . .  | 96         |
| 2.10     | Uninstalling the provision tool . . . . .           | 105        |
| <b>3</b> | <b>Features</b>                                     | <b>107</b> |
| 3.1      | Network . . . . .                                   | 107        |
| 3.2      | Centralized authentication on the cluster . . . . . | 108        |
| 3.3      | Shared and distributed storage deployment . . . . . | 114        |
| 3.4      | GPU accelerator configuration . . . . .             | 118        |
| 3.5      | Offline repositories for the cluster . . . . .      | 118        |
| 3.6      | Additional utilities . . . . .                      | 120        |
| 3.7      | Telemetry and visualizations . . . . .              | 128        |
| <b>4</b> | <b>Logging</b>                                      | <b>153</b> |
| 4.1      | Log management . . . . .                            | 153        |
| 4.2      | Control plane logs . . . . .                        | 154        |
| 4.3      | Logs of individual containers . . . . .             | 156        |
| 4.4      | Provisioning logs . . . . .                         | 156        |
| 4.5      | Telemetry logs . . . . .                            | 157        |
| 4.6      | Grafana Loki . . . . .                              | 157        |
| <b>5</b> | <b>Troubleshooting</b>                              | <b>159</b> |
| 5.1      | Known issues . . . . .                              | 159        |

|           |   |            |
|-----------|---|------------|
| 5.2       | Frequently asked questions . . . . .                          | 168        |
| 5.3       | Troubleshooting guide . . . . .                               | 174        |
| <b>6</b>  | <b>Security Configuration Guide</b>                           | <b>181</b> |
| 6.1       | Preface . . . . .   | 181        |
| 6.2       | Security Quick Reference . . . . .                            | 183        |
| 6.3       | Product and Subsystem Security . . . . .                      | 184        |
| 6.4       | Authentication to external systems . . . . .                  | 186        |
| 6.5       | Network security . . . . .                                    | 186        |
| 6.6       | Miscellaneous Configuration and Management Elements . . . . . | 190        |
| <b>7</b>  | <b>Sample Files</b>   | <b>193</b> |
| 7.1       | inventory file . . . . .                                      | 193        |
| 7.2       | pxe_mapping_file.csv . . . . .                                | 193        |
| 7.3       | switch_inventory . . . . .                                    | 194        |
| 7.4       | powervault_inventory . . . . .                                | 194        |
| 7.5       | NFS Server inventory file . . . . .                           | 194        |
| <b>8</b>  | <b>Limitations</b>  | <b>195</b> |
| <b>9</b>  | <b>Best Practices</b>   | <b>197</b> |
| <b>10</b> | <b>Contributing To Omnia</b>                                  | <b>199</b> |
| 10.1      | Creating A Pull Request . . . . .                             | 199        |



Ansible playbook-based deployment of Slurm and Kubernetes on servers running an RPM-based Linux OS.

**Omnia** (Latin: all or everything) is a deployment tool to turn servers with RPM-based Linux images into functioning Slurm/Kubernetes clusters.

**Omnia** is an open source project hosted on [GitHub](#). Go to [GitHub](#) to view the source, open issues, ask questions, and participate in the project.

### Licensing

Omnia is made available under the [Apache 2.0 license](#).

---

**Note:** Omnia playbooks are licensed under the Apache 2.0 license. Once an end-user initiates Omnia, that end-user will enable deployment of other open source software that is licensed separately by their respective developer communities. For a comprehensive list of software and their licenses, [click here](#). Dell (or any other contributors) shall have no liability regarding and no responsibility to provide support for an end-users use of any open source software and end-users are encouraged to ensure that they are complying with all such licenses. Omnia is provided “as is” without any warranty, express or implied. Dell (or any other contributors) shall have no liability for any direct, indirect, incidental, punitive, special, or consequential damages for an end-users use of Omnia.

---

For a better understanding of what Omnia does, check out our [docs](#)!

### Omnia Community Members



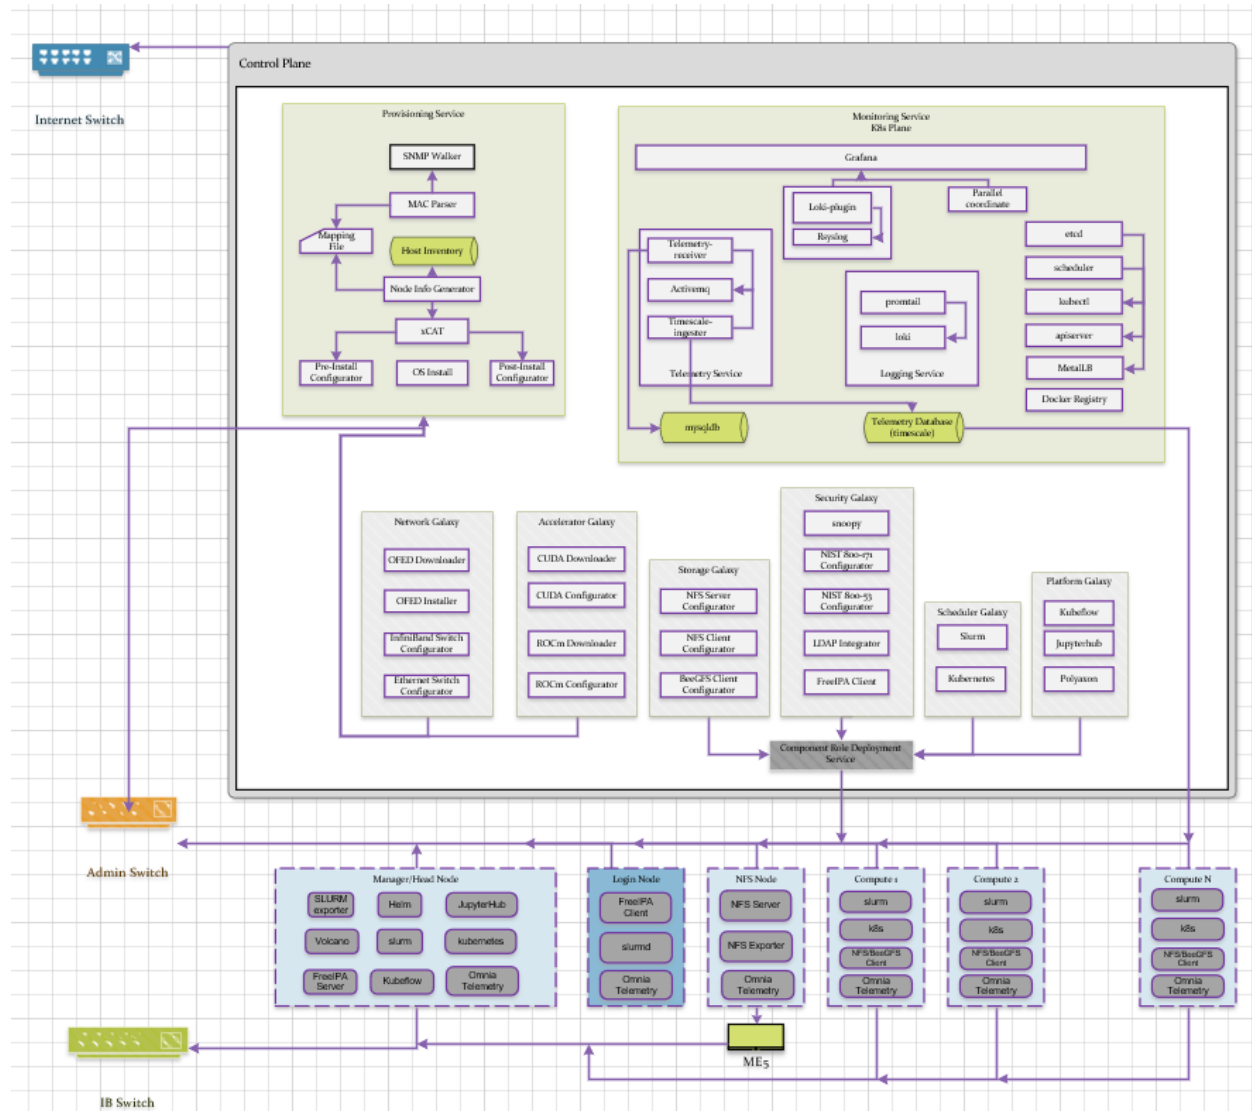
### Table Of Contents



## **OMNIA: OVERVIEW**

**Omnia** (Latin: all or everything) is a deployment tool to configure Dell PowerEdge servers running standard RPM-based Linux OS images into clusters capable of supporting HPC, AI, and data analytics workloads. It uses Slurm, Kubernetes, and other packages to manage jobs and run diverse workloads on the same converged solution. It is a collection of Ansible playbooks, is open source, and is constantly being extended to enable comprehensive workloads.

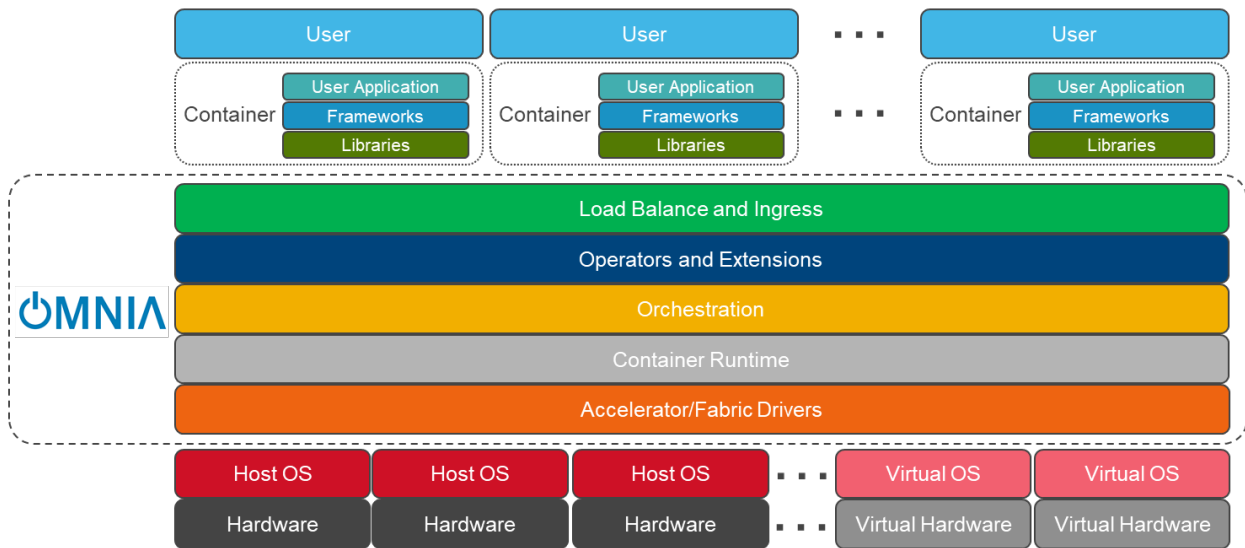
## 1.1 Architecture



### 1.1.1 Omnia stack

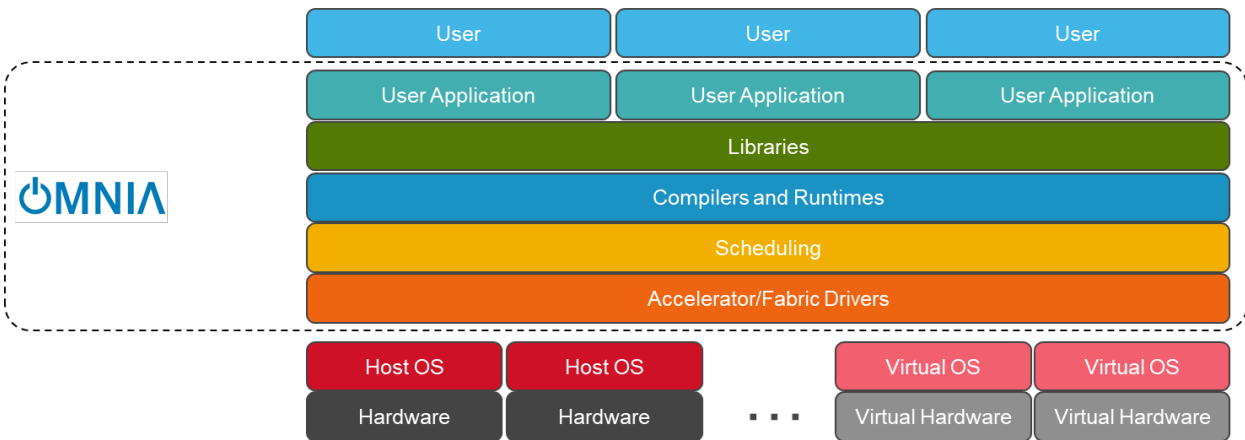
#### Kubernetes

## Omnia Kubernetes stack



### Slurm

## Omnia Slurm stack



If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

## 1.2 New Features

- Expanded [telemetry collection](#) support to Regular, health check and GPU metrics.
- [Rsyslog](#) : Added ability to aggregate logs via xCAT's syslog.
- Integration of apptainer for [containerized HPC benchmark execution](#).
- Integration of [grafana and loki installation](#) within `telemetry.yml`.

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

## 1.3 Releases

### 1.3.1 1.5

- Expanded [telemetry collection](#) support to Regular, health check and GPU metrics.
- [Rsyslog](#) : Added ability to aggregate logs via xCAT's syslog.
- Integration of apptainer for [containerized HPC benchmark execution](#).
- Integration of [grafana and loki installation](#) within `telemetry.yml`.

### 1.3.2 1.4.3

- XE 9640, R760 XA, R760 XD2 are now supported as control planes or target nodes with Nvidia H100 accelerators.
- Added ability for split port configuration on NVIDIA Quantum-2-based QM9700 (Nvidia InfiniBand NDR400 switches).
- Extended password-less SSH support for multiple user configuration in a single execution.
- Input mapping files and inventory files now support commented entries for customized playbook execution.
- NFS share is now available for hosting user home directories within the cluster.

### 1.3.3 1.4.2

- XE9680, R760, R7625, R6615, R7615 are now supported as control planes or target nodes.
- Added ability for switch-based discovery of remote servers and PXE provisioning.
- Active RedHat subscription is no longer required on the control plane and the cluster nodes. Users can configure and use local RHEL repositories.
- IP ranges can be defined for assignment to remote nodes when discovered via the switch.

### 1.3.4 1.4.1

- R660, R6625 and C6620 platforms are now supported as control planes or target nodes.
- One touch provisioning now allows for OFED installation, NVIDIA CUDA-toolkit installation along with iDRAC and InfiniBand IP configuration on target nodes.
- Potential servers can now be discovered via iDRAC.
- Servers can be provisioned automatically without manual intervention for booting/PXE settings.
- Target node provisioning status can now be checked on the control plane by viewing the OmniaDB.
- Omnia clusters can be configured with password-less SSH for seamless execution of HPC jobs run by non-root users.
- Accelerator drivers can be installed on Rocky target nodes in addition to RHEL.

### 1.3.5 1.4

- Provisioning of remote nodes through PXE boot by providing TOR switch IP
- Provisioning of remote nodes through PXE boot by providing mapping file
- PXE provisioning of remote nodes through admin NIC or shared LOM NIC
- Database update of mac address, hostname and admin IP
- Optional monitoring support(Grafana installation) on control plane
- OFED installation on the remote nodes
- CUDA installation on the remote nodes
- AMD accelerator and ROCm support on the remote nodes
- Omnia playbook execution with Kubernetes, Slurm & FreeIPA installation in all cluster nodes
- Infiniband switch configuration and split port functionality
- Added support for Ethernet Z series switches.

### 1.3.6 1.3

- CLI support for all Omnia playbooks (AWX GUI is now optional/deprecated).
- Automated discovery and configuration of all devices (including PowerVault, InfiniBand, and ethernet switches) in shared LOM configuration.
- Job based user access with Slurm.
- AMD server support (R6415, R7415, R7425, R6515, R6525, R7515, R7525, C6525).
- PowerVault ME5 series support (ME5012, ME5024, ME5084).
- PowerVault ME4 and ME5 SAS Controller configuration and NFS server, client configuration.
- NFS bolt-on support.
- BeeGFS bolt-on support.
- Lua and Lmod installation on manager and compute nodes running RedHat 8.x, Rocky 8.x and Leap 15.3.
- Automated setup of FreeIPA client on all nodes.

- Automate configuration of PXE device settings (active NIC) on iDRAC.

### **1.3.7 1.2.2**

- Bugfix patch release to address AWX Inventory not being updated.

### **1.3.8 1.2.1**

- HPC cluster formation using shared LOM network
- Supporting PXE boot on shared LOM network as well as high speed Ethernet or InfiniBand path.
- Support for BOSS Control Card
- Support for RHEL 8.x with ability to activate the subscription
- Ability to upgrade Kernel on RHEL
- Bolt-on Support for BeeGFS

### **1.3.9 1.2.0.1**

- Bugfix patch release which address the broken cobbler container issue.
- Rocky 8.6 Support

### **1.3.10 1.2**

- Omnia supports Rocky 8.5 full OS on the Control Plane
- Omnia supports ansible version 2.12 (ansible-core) with python 3.6 support
- All packages required to enable the HPC/AI cluster are deployed as a pod on control plane
- Omnia now installs Grafana as a single pane of glass to view logs, metrics and telemetry visualization
- cluster node provisioning can be done via PXE and iDRAC
- Omnia supports multiple operating systems on the cluster including support for Rocky 8.5 and OpenSUSE Leap 15.3
- Omnia can deploy cluster nodes with a single NIC.
- All Cluster metrics can be viewed using Grafana on the Control plane (as opposed to checking the manager node on each cluster)
- AWX node inventory now displays service tags with the relevant operating system.
- Omnia adheres to most of the requirements of NIST 800-53 and NIST 800-171 guidelines on the control plane and login node.
- Omnia has extended the FreeIPA feature to provide authentication and authorization on Rocky Nodes.
- Omnia uses [389ds](<https://directory.fedoraproject.org/>) to provide authentication and authorization on Leap Nodes.
- Email Alerts have been added in case of login failures.
- Administrator can restrict users or hosts from accessing the control plane and login node over SSH.



- Malicious or unwanted network software access can be restricted by the administrator.
- Admins can restrict the idle time allowed in an ssh session.
- Omnia installs apparmor to restrict program access on leap nodes.
- Security on audit log access is provided.
- Program execution on the control plane and login node is logged using snoop tool.
- User activity on the control plane and login node is monitored using psacct/acct tools installed by Omnia
- Omnia fetches key performance indicators from iDRACs present in the cluster
- Omnia also supports fetching performance indicators on the nodes in the cluster when SLURM jobs are running.
- The telemetry data is plotted on Grafana to provide better visualization capabilities.
- Four visualization plugins are supported to provide and analyze iDRAC and Slurm data.
  - Parallel Coordinate
  - Spiral
  - Sankey
  - Stream-net (aka. Power Map)
- In addition to the above features, changes have been made to enhance the performance of Omnia.

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

## 1.4 Support Matrix

### 1.4.1 Hardware Supported by Omnia

#### Servers

#### PowerEdge servers

| Server Type | Server Model  |
|-------------|---|
| 14G         | C4140 C6420 R240 R340 R440 R540 R640 R740 R740xd R740xd2 R840 R940 R940xa |
| 15G         | C6520 R650 R750 R750xa  |
| 16G         | C6620 R660 R6625 R760 XE8640 R760xa <sup>1</sup> R760xd2 XE9680           |

**Note:** Since Cloud Enclosures only support shared LOM connectivity, it is recommended that BMC or Switch-based methods of discovery are used.

<sup>1</sup> The R760xa supports both H100 and A100.

## AMD servers

| Server Type | Server Model                  |
|-------------|-------------------------------|
| 14G         | R6415 R7415 R7425             |
| 15G         | R6515 R6525 R7515 R7525 C6525 |
| 16G         | R6625 R7625 R7615 R6615       |

New in version 1.2: 15G servers

New in version 1.3: AMD servers

New in version 1.4.1: Intel 16G servers

New in version 1.4.3: Intel: R760 XE8640 R760xa R760xd2 XE9680; AMD 16G servers

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

## Storage

### Powervault Storage

| Storage Type | Storage Model        |
|--------------|----------------------|
| ME4          | ME4084 ME4024 ME4012 |
| ME5          | ME5012 ME5024 ME5084 |

New in version 1.3: PowerVault ME5 storage support

### BOSS Controller Cards

| BOSS Controller Model | Drive Type  |
|-----------------------|---|
| T2GFX                 | EC, 5300, SSD, 6GBPS SATA, M.2, 512E, ISE, 240GB  |
| M7F5D                 | EC, S4520, SSD, 6GBPS SATA, M.2, 512E, ISE, 480GB |

New in version 1.2.1: BOSS controller cards

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

## Switches

| Switch Type          | Switch Model   |
|----------------------|--|
| Mellanox<br>Switches | InfiniBand<br>NVIDIA MQM8700-HS2F Quantum HDR InfiniBand Switch 40 QSFP56 NVIDIA<br>QUANTUM-2 QM9700 |

| Switch Type      |            | Switch Model           |                       |                       |                        |
|------------------|------------|------------------------|-----------------------|-----------------------|------------------------|
| Dell<br>Switches | Networking | PowerSwitch S3048-ON   | PowerSwitch S5232F-ON | PowerSwitch Z9264F-ON | PowerSwitch N3248TE-ON |
|                  |            | PowerSwitch N3248TE-ON |                       |                       |                        |

**Note:**

- The switches that have reached EOL might not function properly. It is recommended by Omnia to use the switch models mentioned in support matrix.
- Omnia requires that OS10 be installed on ethernet switches.
- Omnia requires that MLNX-OS be installed on Infiniband switches.

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

## 1.4.2 Operating Systems

### Red Hat Enterprise Linux

| OS Version | Control Plane | Cluster Nodes |
|------------|---------------|---------------|
| 8.1        | No            | Yes           |
| 8.2        | No            | Yes           |
| 8.3        | No            | Yes           |
| 8.4        | Yes           | Yes           |
| 8.5        | Yes           | Yes           |
| 8.6        | Yes           | Yes           |
| 8.7        | Yes           | Yes           |

**Note:**

- Always deploy the DVD Edition of the OS on cluster nodes to access offline repos.
- For RHEL 8.5 and below, ensure that RHEL subscription is enabled OR sshpass is available to install or download to the control plane (from any local repository).

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

## Rocky

**Caution:** THE ROCKY LINUX OS VERSION ON THE CLUSTER WILL BE UPGRADED TO THE LATEST 8.x VERSION AVAILABLE IRRESPECTIVE OF THE PROVISION\_OS\_VERSION PROVIDED IN PROVISION\_CONFIG.YML.

| OS Version | Control Plane | Cluster Nodes |
|------------|---------------|---------------|
| 8.4        | Yes           | No            |
| 8.5        | Yes           | No            |
| 8.6        | Yes           | No            |
| 8.7        | Yes           | No            |
| 8.8        | Yes           | Yes           |

---

**Note:**

- Always deploy the DVD (Full) Edition of the OS on cluster Nodes.
  - AMD ROCm driver installation is not supported by Omnia on Rocky cluster nodes.
- 

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

### 1.4.3 Testing matrix

## Hosts

| Generation | Server Model                           |
|------------|--|
| 14G        | PowerEdge C6420                        |
| 14G        | PowerEdge R640                         |
| 14G        | PowerEdge R740xd                       |
| 14G        | PowerEdge R740xc                       |
| 14G        | PowerEdge R540                         |
| 14G        | PowerEdge R440                         |
| 15G        | PowerEdge R550                         |
| 15G        | PowerEdge R650                         |
| 15G        | PowerEdge R450                         |
| 15G        | PowerEdge R6525                        |
| 15G        | PowerEdge R750xa                       |
| 15G        | PowerEdge R6515                        |
| 15G        | PowerEdge R7525                        |
| 16G        | PowerEdge R760                         |
| 16G        | PowerEdge R760xa                       |
| 16G        | PowerEdge R6615                        |
| 16G        | PowerEdge R6625                        |
| 16G        | PowerEdge R7615                        |
| 16G        | PowerEdge R7625                        |
| 16G        | PowerEdge C6620                        |
| 16G        | PowerEdge R660                         |
| 16G        | PowerEdge R760                         |
| 16G        | PowerEdge XE9640 <sup>1</sup>          |
| 16G        | PowerEdge XE9680 <sup>Page 13, 1</sup> |

## NICs

| NIC  |
|--|
| Intel(R) Ethernet 10G 4P X710/I350 rNDC          |
| Intel(R) Ethernet Converged Network Adapter X710 |
| Mellanox ConnectX-5 Single Port 100 GbE QSFP+    |
| Mellanox ConnectX-5 Single Port 0 GbE QSFP       |
| I350GbE Controller                               |
| Broadcom Adv Dual 25Gb Ethernet                  |
| Mellanox ConnectX-6 Single Port VPI HDR QSFP     |
| Mellanox ConnectX-5 Single Port 56 GbE QSFP+     |
| Mellanox ConnectX-6 Single Port VPI HDR 100 QSFP |
| Broadcom Gigabit Ethernet BCM5720                |
| Broadcom Adv Dual 10GBASE-t Ethernet             |
| Mellanox Network Adapter (10 Gb)                 |
| Mellanox ConnectX-5 Ex 100 GbE QSFP              |

<sup>1</sup> This model was tested using Omnia 1.4.3.

## GPUs

| GPU         |
|-------------|
| Nvidia - T4 |
| AMD - MI200 |

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

## 1.4.4 Software Installed by Omnia

| OSS Title                 | License Name/Version #                 | Description                                   |
|---------------------------|--|---|
| Slurm Workload manager    | GNU General Public License             | HPC Workload Manager                          |
| Kubernetes Controllers    | Apache-2.0                             | HPC Workload Manager                          |
| MariaDB                   | GPL 2.0                                | Relational database used by Slurm             |
| Docker CE                 | Apache-2.0                             | Docker Service                                |
| NVidia container runtime  | Apache-2.0                             | Nvidia container runtime library              |
| Python-pip                | MIT License                            | Python Package                                |
| kubelet                   | Apache-2.0                             | Provides external, versioned ComponentCon     |
| kubeadm                   | Apache-2.0                             | “fast paths” for creating Kubernetes clusters |
| kubectl                   | Apache-2.0                             | Command line tool for Kubernetes              |
| jupyterhub                | BSD-3Clause New or Revised License     | Multi-user hub                                |
| kfctl                     | Apache-2.0                             | CLI for deploying and managing Kubeflow       |
| kubeflow                  | Apache-2.0                             | Cloud Native platform for machine learning    |
| helm                      | Apache-2.0                             | Kubernetes Package Manager                    |
| tensorflow                | Apache-2.0                             | Machine Learning framework                    |
| horovod                   | Apache-2.0                             | Distributed deep learning training framework  |
| MPI                       | 3Clause BSD License                    | HPC library                                   |
| spark                     | Apache-2.0                             |   |
| coreDNS                   | Apache-2.0                             | DNS server that chains plugins                |
| cni                       | Apache-2.0                             | Networking for Linux containers               |
| dellemc.openmanage        | GNU-General Public License v3.0        | OpenManage Ansible Modules simplifies an      |
| dellemc.os10              | GNU-General Public License v3.0        | It provides networking hardware abstraction   |
| community.general ansible | GNU-General Public License v3.0        | The collection is a part of the Ansible packa |
| redis                     | BSD-3-Clause License                   | In-memory database                            |
| cri-o                     | Apache-2.0                             | CRI-O is an implementation of the Kuberne     |
| buildah                   | Apache-2.0                             | Tool to build and run containers              |
| OpenSM                    | GNU General Public License 2           |   |
| omsdk                     | Apache-2.0                             | Dell EMC OpenManage Python SDK (OMS           |
| freeipa                   | GNU General Public License v3          | Authentication system used on the login nod   |
| bind-dyndb-ldap           | GNU General Public License v2          | LDAP driver for BIND9. It allows you to re    |
| slurm-exporter            | GNU General Public License v3          | Prometheus collector and exporter for metri   |
| prometheus                | Apache-2.0                             | Open-source monitoring system with a dime     |
| singularity               | BSD License                            | Container platform. It allows you to create a |
| loki                      | GNU AFFERO GENERAL PUBLIC LICENSE v3.0 | Loki is a log aggregation system designed to  |
| promtail                  | Apache-2.0                             | Promtail is an agent which ships the content  |
| Kube prometheus stack     | Apache-2.0                             | Kube Prometheus Stack is a collection of Ku   |
| mailx                     | MIT License                            | mailx is a Unix utility program for sending a |
| xorriso                   | GPL 3.0                                | xorriso copies file objects from POSIX com    |
| openshift                 | Apache-2.0                             | On-premises platform as a service built arou  |

| OSS Title                        | License Name/Version #             | Description   |
|----------------------------------|------------------------------------|---|
| grafana                          | GNU AFFERO GENERAL PUBLIC LICENSE  | Grafana is the open source analytics & monitoring dashboard   |
| kubernetes.core                  | GPL 3.0                            | Performs CRUD operations on K8s objects   |
| community.grafana                | GPL 3.0                            | Technical Support for open source grafana   |
| activemq                         | Apache-2.0                         | Most popular multi protocol, message broker   |
| golang                           | BSD-3-Clause License               | Go is a statically typed, compiled programming language   |
| mysql                            | GPL 2.0                            | MySQL is an open-source relational database management system   |
| postgresSQL                      | PostgresSQL License                | PostgreSQL, also known as Postgres, is a free and open-source relational database management system                                 |
| idrac-telemetry-reference tools  | Apache-2.0                         | Reference toolset for PowerEdge telemetry management  |
| nsfcac/grafana-plugin            | MIT License                        | Machine Learning Framework  |
| jansson                          | MIT License                        | C library for encoding, decoding and manipulating JSON data   |
| libjwt                           | Mozilla Public License-2.0 License | JWT C Library   |
| 389-ds                           | GPL                                | LDAP server used for authentication, access control and directory services  |
| apparmor                         | GNU General Public License         | Controls access based on paths of the program   |
| snoopy                           | GPL 2.0                            | Snoopy is a small library that logs all program activity  |
| timescaledb                      | Apache-2.0                         | TimescaleDB is a time-series SQL database extension for PostgreSQL  |
| Beegfs-Client                    | GPLv2                              | BeeGFS is a high-performance parallel file system   |
| redhat subscription              | Apache-2.0                         | Red Hat Subscription Management (RHSM)  |
| Lmod                             | MIT License                        | Lmod is a Lua based module system that eases the management of software environments  |
| Lua                              | MIT License                        | Lua is a lightweight, high-level, multi-paradigm programming language   |
| ansible posix                    | GNU General Public License         | Ansible Collection targeting POSIX and POSIX-like systems   |
| xCAT                             | Eclipse Public License 1.0         | Provisioning tool that also creates custom distributions  |
| CUDA Toolkit                     | NVIDIA License                     | The NVIDIA® CUDA® Toolkit provides a complete set of tools and libraries for developing and optimizing GPU-accelerated applications |
| MLNX-OFED                        | BSD License                        | MLNX_OFED is an NVIDIA tested and packaged Mellanox OpenFabrics Enterprise Distribution (OFED) driver                               |
| ansible pylibssh                 | LGPL 2.1                           | Python bindings to client functionality of libssh   |
| perl-DBD-Pg                      | GNU General Public License v3      | DBD::Pg - PostgreSQL database driver for Perl   |
| ansible.utils ansible collection | GPL 3.0                            | Ansible Collection with utilities to ease the management of Ansible   |
| pandas                           | BSD-3-Clause License               | pandas is a fast, powerful, flexible and easy-to-use data structure for data analysis   |
| python3-netaddr                  | BSD License                        | A Python library for representing and manipulating IP networks  |
| psycpg2-binary                   | GNU Lesser General Public License  | Psycpg is the most popular PostgreSQL database driver for Python  |
| python.requests                  | Apache-2.0                         | Makes HTTP requests simpler and more human-friendly   |

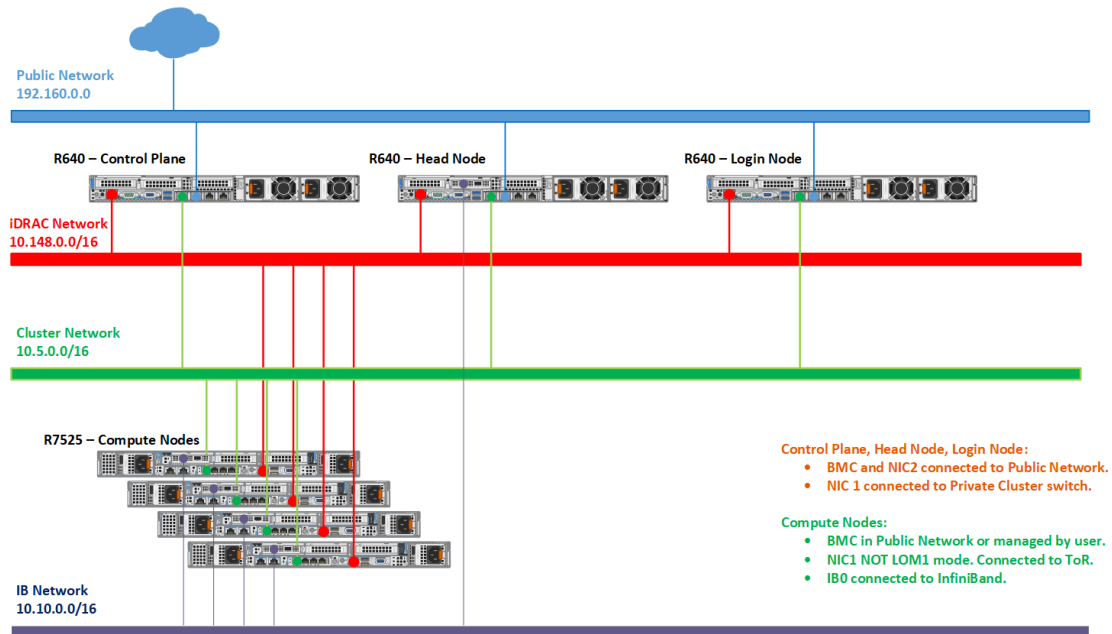
If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

## 1.5 Network Topologies

### 1.5.1 Network Topology: Dedicated Setup

Depending on internet access for host nodes, there are two ways to achieve a dedicated NIC setup:



#### 1. Dedicated Setup with dedicated public NIC on cluster nodes

When all cluster nodes have their own public network access, `primary_dns` and `secondary_dns` in `provision_config.yml` become optional variables as the control plane is not required to be a gateway to the network. The network design would follow the below diagram:

#### 2. Dedicated Setup with single NIC on cluster nodes

When all cluster nodes rely on the control plane for public network access, the variables `primary_dns` and `secondary_dns` in `provision_config.yml` are used to indicate that the control plane is the gateway for all cluster nodes to get internet access. Since all public network traffic will be routed through the control plane, the user may have to take precautions to avoid bottlenecks in such a set-up.

#### Recommended discovery mechanism

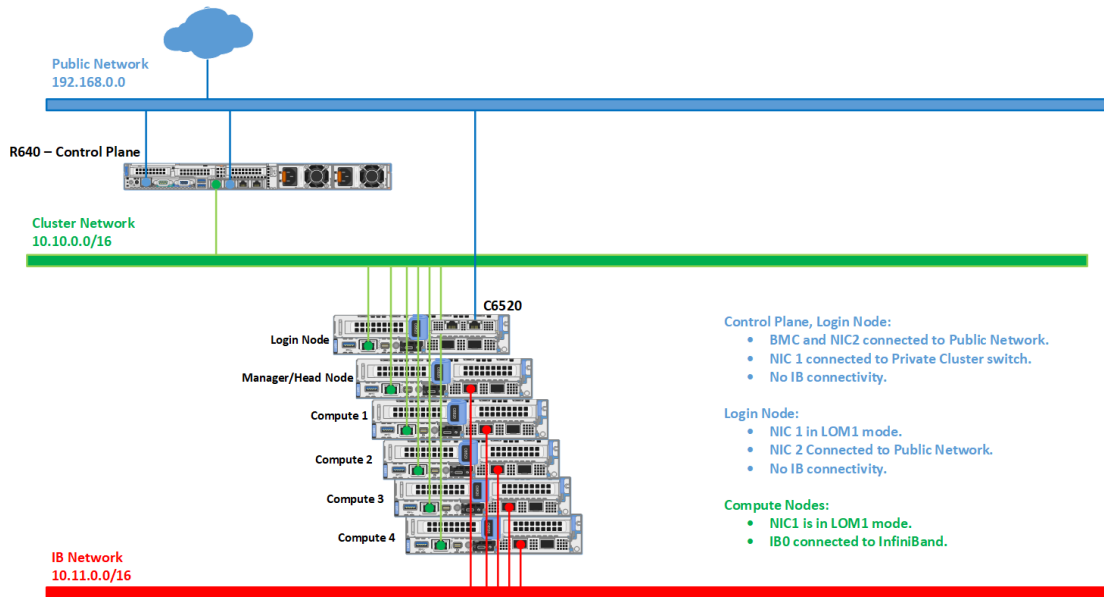
- `mapping`
- `snmpwalk`

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).



## 1.5.2 Network Topology: LOM Setup

A LOM port could be shared with the host operating system production traffic. Also, LOM ports can be dedicated to server management. For example, with a four-port LOM adapter, LOM ports one and two could be used for production data while three and four could be used for iDRAC, VNC, RDP, or other operating system-based management data.



**Caution:** Target nodes using LOM ports cannot be discovered via snmpwalk.


### Recommended discovery mechanism

- mapping
- bmc
- switch-based

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

### 1.5.3 Network Topology: Hybrid setup

For an environment containing both LOM and BMC ports, the provision tool needs to be run twice to correctly manage all servers in the network.



The first time the provision tool is run (to discover the dedicated BMC ports), ensure that the following variables are set in `input/provision_config.yml`:

- `network_interface_type`: dedicated
- `discovery_mechanism`: mapping
- Leave the variables `bmc_nic_subnet`, `bmc_static_start_range` and `bmc_static_end_range` blank in `input/provision_config.yml`. Entering these variables will cause IP reassignment and can interfere with the availability of ports on your target servers.
- Do not use the following methods to discover nodes in a Hybrid setup: `snmpwalk`, `switch-based`.

---

**Note:** For users who don't want internet access routed to target nodes, keep `primary_dns` and `secondary_dns` blank in `input/provision_config.yml`.

---

Once all the dedicated NICs are discovered, re-run the provisioning tool (to discover the shared LOM ports) with the following variables in `input/provision_config.yml`:

- `network_interface_type`: lom
- `discovery_mechanism`: bmc

To assign BMC NICs and route internet access to your target nodes, populate the values of `bmc_nic_subnet`, `bmc_static_start_range`, `bmc_static_end_range`, `primary_dns` and `secondary_dns` in `input/provision_config.yml` during the second run of the provision tool.

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

## 1.6 Blogs about Omnia

- [Introduction to Omnia](#)
- [Taming the Accelerator Cambrian Explosion with Omnia](#)
- [Containerized HPC Workloads Made Easy with Omnia and Singularity](#)
- [Solution Overview: Dell Omnia Software](#)
- [Solution Brief: Omnia Software](#)

## 1.7 What Omnia does

Omnia can deploy and configure devices, and build clusters that use Slurm or Kubernetes (or both) for workload management. Omnia will install software from a variety of sources, including:

- Helm repositories
- Source code repositories

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).



## QUICK INSTALLATION GUIDE

Choose a server outside your intended cluster to function as your control plane.

The control plane needs to be internet-capable with Github and a full OS installed.

---

**Note:** Omnia can be run on control planes running RHEL and Rocky. For a complete list of versions supported, check out the [Support Matrix](#) .

---

```
dnf install git -y
```

---

**Note:** Optionally, if the control plane has an Infiniband NIC installed, run the below command:

```
yum groupinstall "Infiniband Support" -y
```

---

Once the Omnia repository has been cloned on to the control plane:

```
git clone https://github.com/dell/omnia.git
```

Change directory to Omnia using:

```
cd omnia  
sh prereq.sh
```

Run the script `prereq.sh` to verify the system is ready for Omnia deployment.

---

**Note:** The permissions on the Omnia directory are set to **0755** by default. Do not change these values.

---

### 2.1 Running prereq.sh

`prereq.sh` is used to install the software utilized by Omnia on the control plane including Python (3.8), Ansible (2.12.10).

```
cd omnia  
sh prereq.sh
```

---

**Note:**

- If SELinux is not disabled, it will be disabled by the script and the user will be prompted to reboot the control plane.

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

## 2.2 Installing the provision tool

This playbook achieves the following tasks:

- Discovers potential cluster nodes
- Installs Rocky or RHEL on the nodes
- Assigns admin/infiniband IPs with optional DHCP routing
- Creates access to offline repositories
- Configures a docker registry to pull images from the internet and store them locally
- Optionally installs OFED and CUDA

### 2.2.1 Before you run the provision tool

- (Recommended) Run `prereq.sh` to get the system ready to deploy Omnia. Alternatively, ensure that [Ansible 2.12.10](#) and [Python 3.8](#) are installed on the system. SELinux should also be disabled.
- Set the IP address of the control plane with a /16 subnet mask. The control plane NIC connected to remote servers (through the switch) should be configured with two IPs (BMC IP and admin IP) in a shared LOM or hybrid set up. In the case dedicated network topology, a single IP (admin IP) is required.

```

2: eno1: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP group default qlen 1000
    link/ether 18:66:da:89:d4:68 brd ff:ff:ff:ff:ff:ff
    inet 10.5.255.254/16 brd 10.5.255.255 scope global noprefixroute eno1
        valid_lft forever preferred_lft forever
    inet 10.3.255.254/16 brd 10.3.255.255 scope global noprefixroute eno1
        valid_lft forever preferred_lft forever
    inet6 fe80::1a66:daff:fe89:d468/64 scope link noprefixroute
        valid_lft forever preferred_lft forever

```

Fig. 1: Control plane NIC IP configuration in a LOM setup

```

2: eno8303: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP group default qlen 1000
    link/ether b4:45:06:eb:da:4d brd ff:ff:ff:ff:ff:ff
    inet 10.5.255.254/16 brd 10.5.255.255 scope global noprefixroute eno8303
        valid_lft forever preferred_lft forever
    inet6 fe80::b645:6ff:feeb:da4d/64 scope link noprefixroute
        valid_lft forever preferred_lft forever

```

Fig. 2: Control plane NIC IP configuration in a dedicated setup

- Set the hostname of the control plane using the `hostname. domain name` format.

#### Hostname requirements

- The Hostname should not contain the following characters: , (comma), . (period) or \_ (underscore). However, the **domain name** is allowed commas and periods.
- The Hostname cannot start or end with a hyphen (-).

- No upper case characters are allowed in the hostname.
- The hostname cannot start with a number.
- The hostname and the domain name (that is: `hostname00000x.domain.xxx`) cumulatively cannot exceed 64 characters. For example, if the `node_name` provided in `input/provision_config.yml` is 'node', and the `domain_name` provided is 'omnia.test', Omnia will set the hostname of a target cluster node to 'node00001.omnia.test'. Omnia appends 6 digits to the hostname to individually name each target node.

For example, `controlplane.omnia.test` is acceptable.

```
hostnamectl set-hostname controlplane.omnia.test
```

**Note:** The domain name specified for the control plane should be the same as the one specified under `domain_name` in `input/provision_config.yml`.

- To provision the bare metal servers, download one of the following ISOs to the control plane:
  1. Rocky 8
  2. RHEL 8.x

**Caution:** THE ROCKY LINUX OS VERSION ON THE CLUSTER WILL BE UPGRADED TO THE LATEST 8.x VERSION AVAILABLE IRRESPECTIVE OF THE PROVISION\_OS\_VERSION PROVIDED IN PROVISION\_CONFIG.YML.

**Note:** Ensure the ISO provided has downloaded seamlessly (No corruption). Verify the SHA checksum/ download size of the ISO file before provisioning to avoid future failures.

Note the compatibility between cluster OS and control plane OS below:

| Control Plane OS           | cluster Node OS   Compatibility |     |
|----------------------------|---------------------------------|-----|
| RHEL <sup>1</sup>          | RHEL                            | Yes |
| RHEL <sup>Page 23, 1</sup> | Rocky                           | Yes |
| Rocky                      | Rocky                           | Yes |

- To **optionally** set up CUDA and OFED using the provisioning tool, download the required repositories to the control plane from here to deploy on the target nodes:
  1. For **NVIDIA GPUs**:: CUDA is a parallel computing platform and application programming interface that allows software to use certain types of graphics processing units for general purpose processing, an approach called general-purpose computing on GPUs.
  2. For **Mellanox**: OFED (OpenFabrics Enterprise Distribution) is open-source software for RDMA and kernel bypass applications. OFED can be used in business, research and scientific environments that require highly efficient networks, storage connectivity and parallel computing.

<sup>1</sup> Ensure that control planes running RHEL have an active subscription or are configured to access local repositories. The following repositories should be enabled on the control plane: **AppStream**, **Code Ready Builder (CRB)**, **BaseOS**. For RHEL control planes running 8.5 and below, ensure that `sshpas` is additionally available to install or download to the control plane (from any local repository).

- **Ensure that all connection names under the network manager match their corresponding device names.**

To verify network connection names:

```
nmcli connection
```

To verify the device name:

```
ip link show
```

In the event of a mismatch, edit the file `/etc/sysconfig/network-scripts/ifcfg-<nic name>` using vi editor.

- When discovering nodes via snmpwalk or a mapping file, all target nodes should be set up in PXE mode before running the playbook.
- Nodes provisioned using the Omnia provision tool do not require a RedHat subscription to run `provision.yml` on RHEL target nodes.
- For RHEL target nodes not provisioned by Omnia, ensure that RedHat subscription is enabled on all target nodes. Every target node will require a RedHat subscription.
- Users should also ensure that all repos (AppStream, BaseOS and CRB) are available on the RHEL control plane.

---

**Note:**

- Enable a repository from your RHEL subscription, run the following commands:

```
subscription-manager repos --enable=codeready-builder-for-rhel-8-x86_64-rpms
subscription-manager repos --enable=rhel-8-for-x86_64-appstream-rpms
subscription-manager repos --enable=rhel-8-for-x86_64-baseos-rpms
```

- Enable an offline repository by creating a `.repo` file in `/etc/yum.repos.d/`. Refer the below sample content:

```
[RHEL-8-appstream]

name=Red Hat AppStream repo

baseurl=http://xx.yy.zz/pub/Distros/RedHat/RHEL8/8.6/AppStream/x86_64/os/

enabled=1

gpgcheck=0

[RHEL-8-baseos]

name=Red Hat BaseOS repo

baseurl=http://xx.yy.zz/pub/Distros/RedHat/RHEL8/8.6/BaseOS/x86_64/os/

enabled=1

gpgcheck=0

[RHEL-8-crb]

name=Red Hat CRB repo

baseurl=http://xx.yy.zz/pub/Distros/RedHat/RHEL8/8.6/CRB/x86_64/os/
```

(continues on next page)



(continued from previous page)

```
enabled=1
gpgcheck=0
```

- Verify your changes by running:

```
yum repolist enabled
Updating Subscription Management repositories.
Unable to read consumer identity
This system is not registered with an entitlement server. You can use subscription-
↪manager to register.
    repo id                                repo name
    RHEL-8-appstream-partners              Red Hat_
↪Enterprise Linux 8.6.0 Partners (AppStream)
    RHEL-8-baseos-partners                 Red Hat_
↪Enterprise Linux 8.6.0 Partners (BaseOS)
    RHEL-8-crb-partners                    Red Hat_
↪Enterprise Linux 8.6.0 Partners (CRB)
```

- Uninstall epel-release if installed on the control plane as Omnia configures epel-release on the control plane. To uninstall epel-release, use the following commands:

```
dnf remove epel-release -y
```

- Ensure that the pxe\_nic and public\_nic are in the firewall zone: public.

---

**Note:**

- After configuration and installation of the cluster, changing the control plane is not supported. If you need to change the control plane, you must redeploy the entire cluster.
  - For servers with an existing OS being discovered via BMC, ensure that the first PXE device on target nodes should be the designated active NIC for PXE booting.
- 

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

## 2.2.2 Discovery Mechanisms

Typically, the choice of discovery mechanism depends on the [Network Topology](#) in your setup. Depending on the value of `discovery_mechanism` in `input/provision_config.yml`, potential target servers can be discovered one of four ways:

### mapping

Manually collect PXE NIC information for target servers and manually define them to Omnia using a mapping file using the below format:

#### `pxe_mapping_file.csv`

```
MAC,Hostname,IP
xx:yy:zz:aa:bb:cc,server,10.5.0.101
aa:bb:cc:dd:ee:ff,server2, 10.5.0.102
```

---

#### Note:

- All the file contents mentioned above are case sensitive. The casing of words like `[manager]`, `[compute]`, etc should be consistent with the samples below when creating inventory or mapping files.
  - The hostnames listed in the above file should be exclusively lower-case with no special characters.
  - The MAC address provided in `pxe_mapping_file.csv` should refer to the PXE NIC on the target nodes.
- 

#### Caution:

- Do not remove or comment any lines in the `input/provision_config.yml` file.
- **THE ROCKY LINUX OS VERSION ON THE CLUSTER WILL BE UPGRADED TO THE LATEST 8.x VERSION AVAILABLE IRRESPECTIVE OF THE PROVISION\_OS\_VERSION PROVIDED IN PROVISION\_CONFIG.YML.**
- `admin_nic_subnet`, `ib_nic_subnet` and `bmc_nic_subnet` should have the same subnet mask (Omnia only supports /16 subnet masks currently).

The following parameters need to be populated in `input/provision_config.yml` to discover target nodes using a mapping file.

Table 1: Parameters

| Parameter   | Details  |
|---|--|
| <b>network_interface_type</b><br>string<br>Required | The network type used on the Omnia cluster.<br>Choices: <ul style="list-style-type: none"><li>dedicated &lt;- default</li><li>lom</li></ul>  |
| <b>discovery_mechanism</b><br>string<br>Required    | The mechanism through which Omnia will discover nodes for provisioning. For more information on how the mechanisms work, go to <a href="#">DiscoveryMechanisms</a> .<br>Choices: <ul style="list-style-type: none"><li>switch_based</li><li>mapping &lt;-default</li><li>bmc</li><li>snmpwalk</li></ul>                |
| <b>provision_os</b><br>string<br>Required           | The operating system to be provisioned on target nodes in the cluster.<br>Choices: <ul style="list-style-type: none"><li>rhel &lt;-default</li><li>rocky</li></ul> <div>Caution:<br/>THE ROCKY LINUX OS VERSION ON THE CLUSTER WILL BE UP-GRADED TO THE LATEST 8.x VERSION AVAILABLE IRRESPECTIVE OF THE PRO-VI-</div> |

**Caution:** The IP address *192.168.25.x* is used for PowerVault Storage communications. Therefore, do not use this IP address for other configurations.

---

**Note:** The input/provision\_config.yml file is encrypted on the first run of the provision tool:

To view the encrypted parameters:

```
ansible-vault view provision_config.yml --vault-password-file .provision_vault_  
↪key
```

To edit the encrypted parameters:

```
ansible-vault edit provision_config.yml --vault-password-file .provision_vault_  
↪key
```

---

To continue to the next steps:

- [Provisioning the cluster](#)

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

## switch\_based

### Pre requisites

- IP address for ToR switch needs to be provided.
- Switch port range where all BMC NICs are connected should be provided.
- SNMP v3 should be enabled on the switch.
- Non-admin user credentials for the switch need to be provided.

---

### Note:

- **To create an SNMPv3 user on S series switches (running OS10), use the following commands:**
  - To create SNMP view: `snmp-server view test_view internet included`
  - To create SNMP group: `snmp-server group testgroup 3 auth read test_view`
  - To create SNMP users: `snmp-server user authuser1 testgroup 3 auth sha authpasswd1`
- **To verify the changes made, use the following commands:**
  - To view the SNMP views: `show snmp view`
  - To view the SNMP groups: `show snmp group`
  - To view the SNMP users: `show snmp user`
- To save this configuration for later use, run: `copy running-configuration startup-configuration`
- For more information on SNMP on S series switch [click here](#)
- For more information on SNMP on N series switch [click here](#)

---

<sup>1</sup> Boolean parameters do not need to be passed with double or single quotes.

- IPMI over LAN needs to be enabled for the BMC.

```
racadm set iDRAC.IPMILan.Enable 1
racadm get iDRAC.IPMILan
```

- BMC credentials should be the same across all servers and provided as input to Omnia.
- Target servers should be configured to boot in PXE mode with appropriate NIC as the first boot device.
- Set the IP address of the control plane with a /16 subnet mask. The control plane NIC connected to remote servers (through the switch) should be configured with two IPs (BMC IP and admin IP) in a shared LOM or hybrid set up. In the case dedicated network topology, a single IP (admin IP) is required.

```
2: eno1: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP group default qlen 1000
    link/ether 18:66:da:89:d4:68 brd ff:ff:ff:ff:ff:ff
    inet 10.5.255.254/16 brd 10.5.255.255 scope global noprefixroute eno1
        valid_lft forever preferred_lft forever
    inet 10.3.255.254/16 brd 10.3.255.255 scope global noprefixroute eno1
        valid_lft forever preferred_lft forever
    inet6 fe80::1a66:daff:fe89:d468/64 scope link noprefixroute
        valid_lft forever preferred_lft forever
```

#### Caution:

- Do not use daisy chain ports or the port used to connect to the control plane in `switch_based_details` in `input/provision_config.yml`. This can cause IP conflicts on servers attached to potential target ports.
- Omnia does not validate SNMP switch credentials, if the provision tool is run with incorrect credentials, use the clean-up script and re-run the provision tool with the correct credentials.
- If you are re-provisioning your cluster (that is, re-running the `provision.yml` playbook) after a `clean-up`, ensure to use a different `admin_nic_subnet` in `input/provision_config.yml` to avoid a conflict with newly assigned servers. Alternatively, disable any OS available in the Boot Option Enable/Disable section of your BIOS settings (BIOS Settings > Boot Settings > UEFI Boot Settings) on all target nodes.
- `admin_nic_subnet`, `ib_nic_subnet` and `bmc_nic_subnet` should have the same subnet mask (Omnia only supports /16 subnet masks currently).

#### Note:

- The IP range `x.y.246.1 - x.y.255.253` (where `x` and `y` are provided by the first two octets of `bmc_nic_subnet`) are reserved by Omnia.
- **If any of the target nodes have a pre-provisioned IP, do not use a `bmc_subnet` and/or `ip_start_range/ip_end_range` that encapsulates the pre-provisioned IP.**
  - For example, if there are target nodes hosted at 10.3.0.11 and 10.3.0.12, `bmc_subnet = 10.3.0.0` with `ip_start_range = 10.3.0.1/ ip_end_range = 10.3.0.255` will cause a conflict with newly assigned servers however, `bmc_subnet = 10.3.0.0` with `ip_start_range = 10.3.0.100/ ip_end_range = 10.3.0.150` would be accepted. Alternatively, a different subnet would be acceptable, ie `bmc_subnet = 10.13.0.0`.

The following parameters need to be populated in `input/provision_config.yml` to discover target nodes using a mapping file.

**Caution:**

- Do not remove or comment any lines in the `input/provision_config.yml` file.
- **THE ROCKY LINUX OS VERSION ON THE CLUSTER WILL BE UPGRADED TO THE LATEST 8.x VERSION AVAILABLE IRRESPECTIVE OF THE PROVISION\_OS\_VERSION PROVIDED IN PROVISION\_CONFIG.YML.**

Table 2: Parameters

| Parameter   | Details  |
|---|--|
| <b>network_interface_type</b><br>string<br>Required | The network type used on the Omnia cluster.<br>Choices: <ul style="list-style-type: none"> <li>• <code>dedicated</code> &lt;- default</li> <li>• <code>lom</code></li> </ul>   |
| <b>discovery_mechanism</b><br>string<br>Required    | The mechanism through which Omnia will discover nodes for provisioning. For more information on how the mechanisms work, go to <a href="#">DiscoveryMechanisms</a> .<br>Choices: <ul style="list-style-type: none"> <li>• <code>switch_based</code></li> <li>• <code>mapping</code> &lt;-default</li> <li>• <code>bmc</code></li> <li>• <code>snmpwalk</code></li> </ul>   |
| <b>provision_os</b><br>string<br>Required           | The operating system to be provisioned on target nodes in the cluster.<br>Choices: <ul style="list-style-type: none"> <li>• <code>rhel</code> &lt;-default</li> <li>• <code>rocky</code></li> </ul>  |
| <b>provision_os_version</b><br>string<br>Required   | OS version of <code>provision_os</code> to be installed.<br>Choices: <ul style="list-style-type: none"> <li>• <code>8.0</code></li> <li>• <code>8.1</code></li> <li>• <code>8.2</code></li> <li>• <code>8.3</code></li> <li>• <code>8.4</code></li> <li>• <code>8.5</code></li> <li>• <code>8.6</code> &lt;- default</li> <li>• <code>8.7</code></li> </ul>                |
| <b>iso_file_path</b><br>string<br>Required          | Path where user has placed the iso image that needs to be provisioned on target nodes. Accepted files are Rocky8-DVD or RHEL-8.x-DVD (full OS). <code>iso_file_path</code> should contain the <code>provision_os</code> and <code>provision_os_version</code> values in the filename.<br><b>Default values:</b> <code>"/home/RHEL-8.6.0-20220420.3-x86_64-dvd1.iso"</code> |

continues on next page

Table 2 – continued from previous page

| Parameter  | Details   |
|--|---|
| <b>timezone</b><br>string<br>Required            | Timezone to be used during OS provisioning. Available timezones are provided <a href="#">here</a> .<br>Choices: <ul style="list-style-type: none"><li>• GMT &lt;- default</li><li>• EST</li><li>• CET</li><li>• MST</li><li>• CST6CDT</li><li>• PST8PDT</li></ul>   |
| <b>language</b><br>string<br>Required            | Language to be used during OS provisioning.<br><b>Default values:</b> en-US   |
| <b>default_lease_time</b><br>integer<br>Required | Default lease time for IPs assigned by DHCP. Range: 21600-86400<br><b>Default values:</b> 86400   |
| <b>provision_password</b><br>string<br>Required  | <ul style="list-style-type: none"><li>• Password set for the root account of target nodes during provisioning.</li><li>• Length &gt;= 8 characters</li><li>• Password must not contain -, , ‘, ’</li><li>• The first character of the string should be an alphabet.</li></ul>   |
| <b>postgresdb_password</b><br>string<br>Required | <ul style="list-style-type: none"><li>• Password set for the postgresDB on target nodes during provisioning.</li><li>• Length &gt;= 8 characters</li><li>• Password must not contain -, , ‘, ’</li><li>• The first character of the string should be an alphabet.</li></ul>   |
| <b>node_name</b><br>string<br>Required           | <ul style="list-style-type: none"><li>• Prefix for target node names, if dynamically allocated.</li><li>• Hostname = node_name + ‘0000x’ + domain_name</li><li>• Hostname &lt;= 65 characters</li><li>• Example: servernode00001.Omnia.test , where node_name =servernode, domain_name =Omnia.test , 00001 used by Omnia.</li></ul> <b>Default values:</b> node |

continues on next page

Table 2 – continued from previous page

| Parameter  | Details  |
|--|--|
| <b>domain_name</b><br>string<br>Required             | <ul style="list-style-type: none"> <li>Domain name the user intends to configure on the cluster.</li> <li>Hostname = node_name + '0000x' + domain_name</li> <li>Hostname &lt;= 65 characters</li> <li>Please provide a valid domain name according to the domain name standards.</li> <li>Example: servernode00001.Omnia.test, where node_name=servernode, domain_name=Omnia.test, 00001 used by Omnia.</li> </ul>   |
| <b>public_nic</b><br>string<br>Required              | The nic/ethernet card that is connected to the public internet.<br><b>Default values:</b> eno2   |
| <b>admin_nic</b><br>string<br>Required               | <ul style="list-style-type: none"> <li>Admin NIC of Control Plane. This is the shared LOM NIC.</li> <li>The value of this variable cannot be changed after successfully running provision.yml.</li> </ul> <b>Default values:</b> eno1  |
| <b>admin_nic_subnet</b><br>string<br>Required        | The subnet within which all Admin IPs are assigned.<br><b>Default values:</b> 10.5.0.0   |
| <b>switch_based_details</b><br>JSON List<br>Required | <ul style="list-style-type: none"> <li>JSON list of switches to query for target nodes.</li> <li>Split port ranges are not accepted here. (Ex: 10:5-10:10 will not be valid).</li> <li>Example: <ul style="list-style-type: none"> <li>- { ip: 172.96.28.12, ports: '1-48,49:3,50' }</li> <li>• Example with 2 switches: <ul style="list-style-type: none"> <li>- { ip: 172.96.28.12, ports: '1-48,49:3,50' }</li> <li>- { ip: 172.96.28.14, ports: '1,2,3,5' }</li> </ul> </li> </ul> </li> </ul> |
| <b>switch_snmp3_username</b><br>string<br>Required   | The non-admin SNMPv3 username for the switch.  |
| <b>switch_snmp3_password</b><br>string<br>Required   | <ul style="list-style-type: none"> <li>The non-admin SNMPv3 password for the switch.</li> <li>The first character of the string should be an alphabet.</li> </ul>  |
| <b>ip_start_range</b><br>string<br>Required          | <ul style="list-style-type: none"> <li>The IP start range for all the NICs within the subnets mentioned above.(Admin, BMC, IB).</li> <li>Example: For the range x.y.0.1 to x.y.0.100 with admin_nic_subnet = 10.5.0.0, target nodes will be assigned admin IPs between 10.5.0.1 and 10.5.0.100.</li> </ul>   |

continues on next page



Table 2 – continued from previous page

| Parameter                                   | Details   |
|---|---|
| <b>ip_end_range</b><br>string<br>Required   | <ul style="list-style-type: none"> <li>The IP end range for all the NICs within the subnets mentioned above.(Admin, BMC, IB).</li> <li>Example: For the range x.y.0.1 to x.y.0.100 with admin_nic_subnet = 10.5.0.0, target nodes will be assigned admin IPs between 10.5.0.1 and 10.5.0.100.</li> </ul>  |
| <b>ib_nic_subnet</b><br>string<br>Optional  | <ul style="list-style-type: none"> <li>If provided, Omnia will handle and assign static IPs to cluster node's IB network.</li> <li>Only the last 16 bits/2 octets of IPv4 are dynamic</li> <li>If provided, the DB entry will be in parallel with the pxe_subnet.</li> <li>Example: If admin_ip: 10.5.0.50 and ib_nic_subnet: 10.10.0.0, then ib_ip: 10.10.0.50</li> </ul>  |
| <b>bmc_nic_subnet</b><br>string<br>Required | <ul style="list-style-type: none"> <li>If provided, Omnia will assign static IPs to BMC NICs on the cluster nodes within the provided subnet.</li> <li>If network_interface_type: lom, mandatory for discovery_mechanism: mapping, switch_based and bmc.</li> <li>If network_interface_type: dedicated, optional for discovery_mechanism: mapping, snmp-walk.</li> <li>Note that since the last 16 bits/2 octets of IPv4 are dynamic, please ensure that the parameter value is set to x.x.0.0.</li> <li>When the PXE range and BMC subnet are provided, corresponding NICs will be assigned IPs with the same 3rd and 4th octets.</li> </ul> |
| <b>bmc_username</b><br>string<br>Required   | <ul style="list-style-type: none"> <li>The username for iDRAC.</li> <li>The username must not contain -, ' ,"</li> </ul>  |
| <b>bmc_password</b><br>string<br>Required   | <ul style="list-style-type: none"> <li>The password for iDRAC.</li> <li>The password must not contain -, ' ,"</li> </ul>  |

continues on next page

Table 2 – continued from previous page

| Parameter  | Details   |
|--|---|
| <b>update_repos</b><br>boolean <sup>1</sup><br>Required                            | <ul style="list-style-type: none"> <li>Indicates whether <code>provision.yml</code> will update offline RHEL repos.</li> <li>If <code>update_repos: false</code>, the update repos for BaseOS and AppStream will not be updated to the latest versions available.</li> <li>If <code>update_repos: true</code>, the update repos for BaseOS and AppStream will be updated to the latest versions available.</li> </ul> <hr/> <p><b>Note:</b> By default, AppStream and BaseOS repos will be configured from the given ISO file.</p> <hr/> <p>Choices:<br/> false &lt;- Default<br/> true</p>   |
| <b>rhel_repo_alphabetical_folders</b><br>boolean <sup>Page 36, 1</sup><br>Required | <ul style="list-style-type: none"> <li>Indicates whether the packages in local or subscription repos should be ordered in alphabetical directories.</li> <li>This variable should be filled if control plane OS is RHEL and local RHEL repository is available.</li> </ul> <p>Choices:<br/> false &lt;- Default<br/> true</p>   |
| <b>rhel_repo_local_path</b><br>JSON list<br>Optional                               | <ul style="list-style-type: none"> <li>The repo path and names of the software repository to be configured on the cluster nodes.</li> <li>Provide the repo data file path, which ends with <code>.repo</code> extension in <code>repo_url</code> parameter.</li> <li>Provide a <b>valid</b> url for BaseOS and AppStream repositories. Omnia does not validate the <code>repo_url</code> provided. Invalid entries will cause <code>provision.yml</code> to fail.</li> <li>This variable should be filled if control plane OS is RHEL and subscription is not activated.</li> <li>This variable should be filled if the control plane OS is Rocky and the <code>provision_os</code> is <code>rhel</code>.</li> </ul> <p>Sample value:</p> <pre>- { repo: "AppStream", repo_url: "http://xx.yy.zz/pub/Distros/RedHat/RHEL8/8.8/RHEL-8-appstream.repo",   repo_name: "RHEL-8-appstream-partners" } - { repo: "BaseOS", repo_url: "http://xx.yy.zz/pub/Distros/RedHat/RHEL8/8.8/RHEL-8-baseos.repo",   repo_name: "RHEL-8-baseos-partners" }</pre> |

continues on next page

Table 2 – continued from previous page

| Parameter                                      | Details  |
|--|--|
| <b>primary_dns</b><br>string<br>Optional       | <ul style="list-style-type: none"> <li>The primary DNS host IP queried to provide Internet access to cluster Node (through DHCP routing).</li> <li>Currently, the <code>primary_dns</code> value stored in <code>input/provision_config.yml</code> cannot be part of any of the subnets (<code>admin_nic_subnet</code>, <code>ib_nic_subnet</code> and <code>bmc_nic_subnet</code>) also defined in <code>input/provision_config.yml</code>.</li> </ul> <p>Ex: If the <code>primary_dns</code> is set to 10.15.0.7, the subnet 10.15.0.0 cannot be used for <code>admin_nic_subnet</code>, <code>ib_nic_subnet</code> or <code>bmc_nic_subnet</code>.</p>  |
| <b>secondary_dns</b><br>string<br>Optional     | The secondary DNS host IP queried to provide Internet access to cluster Node (through DHCP routing)  |
| <b>disk_partition</b><br>JSON list<br>Optional | <ul style="list-style-type: none"> <li>User defined disk partition applied to remote servers.</li> <li>The disk partition <code>desired_capacity</code> has to be provided in MB.</li> <li>Valid <code>mount_point</code> values accepted for disk partition are <code>/home</code>, <code>/var</code>, <code>/tmp</code>, <code>/usr</code>, <code>swap</code>.</li> <li>Default partition size provided for <code>/boot</code> is 1024MB, <code>/boot/efi</code> is 256MB and the remaining space to / partition.</li> <li>Values are accepted in the form of JSON list such as: <code>, - { mount_point: "/home", desired_capacity: "102400" }</code></li> </ul> <p><b>Default values:</b> <code>- { mount_point: "", desired_capacity: "" }</code></p> |
| <b>mlnx_ofed_path</b><br>string<br>Optional    | Absolute path to a local copy of the .iso file containing Mellanox OFED packages. The image can be downloaded from <a href="https://network.nvidia.com/products/infiniband-drivers/linux/mlnx_ofed/">https://network.nvidia.com/products/infiniband-drivers/linux/mlnx_ofed/</a> . Sample value: <code>/root/MLNX_OFED_LINUX-5.8-1.1.2.1-rhel8.6-x86_64.iso</code>   |
| <b>cuda_toolkit_path</b><br>string<br>Optional | Absolute path to local copy of .rpm file containing CUDA packages. The cuda rpm can be downloaded from <a href="https://developer.nvidia.com/cuda-downloads">https://developer.nvidia.com/cuda-downloads</a> . CUDA will be installed post provisioning without any user intervention. Eg: <code>cuda_toolkit_path: "/root/cuda-repo-rhel8-12-0-local-12.0.0_525.60.13-1.x86_64.rpm"</code>  |

continues on next page

Table 2 – continued from previous page

| Parameter   | Details   |
|---|---|
| <b>apptainer_support</b><br>boolean <sup>Page 36, 1</sup><br>Required | <ul style="list-style-type: none"> <li>Indicates whether apptainer will be installed on the cluster to enable execution of HPC benchmarks in a containerized environment.</li> <li>If <code>apptainer_support</code>: false, apptainer will not be installed on the cluster.</li> <li>If <code>apptainer_support</code>: true, apptainer will be installed on the cluster.</li> </ul> |

**Note:** The `input/provision_config.yml` file is encrypted on the first run of the provision tool:

To view the encrypted parameters:

```
ansible-vault view provision_config.yml --vault-password-file .provision_vault_
↪key
```

To edit the encrypted parameters:

```
ansible-vault edit provision_config.yml --vault-password-file .provision_vault_
↪key
```

To clear the configuration on Omnia provisioned switches and ports, [click here](#).

To continue to the next steps:

- [Provisioning the cluster](#)

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

## bmc

For automatic provisioning of servers and discovery, the BMC method can be used.

### Pre requisites

- Set the IP address of the control plane with a /16 subnet mask. The control plane NIC connected to remote servers (through the switch) should be configured with two IPs (BMC IP and admin IP) in a shared LOM or hybrid set up. In the case dedicated network topology, a single IP (admin IP) is required.

```
2: eno1: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP group default qlen 1000
    link/ether 18:66:da:89:d4:68 brd ff:ff:ff:ff:ff:ff
    inet 10.5.255.254/16 brd 10.5.255.255 scope global noprefixroute eno1
        valid_lft forever preferred_lft forever
    inet 10.3.255.254/16 brd 10.3.255.255 scope global noprefixroute eno1
        valid_lft forever preferred_lft forever
    inet6 fe80::1a66:daff:fe89:d468/64 scope link noprefixroute
        valid_lft forever preferred_lft forever
```

- BMC NICs should have a static IP assigned or be configured in DHCP mode.
- BMC credentials should be the same across all servers and provided as input to Omnia in the parameters explained below.

<sup>1</sup> Boolean parameters do not need to be passed with double or single quotes.

- Target servers should be configured to boot in PXE mode with the appropriate NIC as the first boot device.
- IPMI over LAN needs to be enabled for the BMC.

```
racadm set iDRAC.IPMLan.Enable 1
racadm get iDRAC.IPMLan
```

- IP ranges (bmc\_static\_start\_range, bmc\_static\_end\_range) provided to Omnia for BMC discovery should be within the same subnet.

#### Caution:

- To create a meaningful range of discovery, ensure that the last two octets of bmc\_static\_end\_range are equal to or greater than the last two octets of the bmc\_static\_start\_range. That is, for the range a.b.c.d - a.b.e.f, e and f should be greater than or equal to c and d. *Ex: 172.20.0.50 - 172.20.1.101 is a valid range however, 172.20.0.101 - 172.20.1.50 is not.*
- If you are re-provisioning your cluster (that is, re-running the provision.yml playbook) after a **clean-up**, ensure to use a different admin\_nic\_subnet in input/provision\_config.yml to avoid a conflict with newly assigned servers. Alternatively, disable any OS available in the Boot Option Enable/Disable section of your BIOS settings (BIOS Settings > Boot Settings > UEFI Boot Settings) on all target nodes.
- admin\_nic\_subnet, ib\_nic\_subnet and bmc\_nic\_subnet should have the same subnet mask (Omnia only supports /16 subnet masks currently).

- All iDRACs should be reachable from the admin\_nic.

#### Note:

##### When iDRACs are in DHCP mode\*

- The IP range *x.y.246.1 - x.y.255.253* (where x and y are provided by the first two octets of bmc\_nic\_subnet) are reserved by Omnia.
- *x.y.246.1 - x.y.250.253* will be the range of IPs reserved for dynamic assignment by Omnia.
- During provisioning, Omnia updates servers to static mode and assigns IPs from *x.y.251.1 - x.y.255.253*.
- Users can see the IPs (that have been assigned from *x.y.251.1 - x.y.255.253*) in the DB after provisioning the servers.
- **For example:**  
If the provided bmc\_subnet is **10.3.0.0** and there are two iDRACs in DHCP mode, the IPs assigned will be **10.3.251.1** and **10.3.251.2**.

The following parameters need to be populated in input/provision\_config.yml to discover target nodes using BMC.

#### Caution:

- Do not remove or comment any lines in the input/provision\_config.yml file.
- **THE ROCKY LINUX OS VERSION ON THE CLUSTER WILL BE UPGRADED TO THE LATEST 8.x VERSION AVAILABLE IRRESPECTIVE OF THE PROVISION\_OS\_VERSION PROVIDED IN PROVISION\_CONFIG.YML.**

Table 3: Parameters

| Parameter   | Details   |
|---|---|
| <b>network_interface_type</b><br>string<br>Required | The network type used on the Omnia cluster.<br>Choices: <ul style="list-style-type: none"><li>• dedicated &lt;- default</li><li>• lom</li></ul>   |
| <b>discovery_mechanism</b><br>string<br>Required    | The mechanism through which Omnia will discover nodes for provisioning. For more information on how the mechanisms work, go to <a href="#">DiscoveryMechanisms</a> .<br>Choices: <ul style="list-style-type: none"><li>• switch_based</li><li>• mapping &lt;-default</li><li>• bmc</li><li>• snmpwalk</li></ul> |

---

continues on next page

Table 3 – continued from previous page

| Parameter                                 | Details  |
|---|--|
| <b>provision_os</b><br>string<br>Required | <p>The operating system to be provisioned on target nodes in the cluster.</p> <p>Choices:</p> <ul style="list-style-type: none"><li>• rhel &lt;-default</li><li>• rocky</li></ul> <div><b>Caution:</b><br/>THE ROCKY LINUX OS VERSION ON THE CLUSTER WILL BE UP-GRADED TO THE LATEST 8.x VERSION AVAILABLE IRRESPECTIVE OF THE PROVIDED IN PROVISION_OS_VERSION IN PROVISION_CONFIG.YML.</div> |

continues on next page

Table 3 – continued from previous page

| Parameter   | Details  |
|---|--|
| <b>provision_os_version</b><br>string<br>Required | OS version of provision_os to be installed.<br>Choices: <ul style="list-style-type: none"><li>• 8.0</li><li>• 8.1</li><li>• 8.2</li><li>• 8.3</li><li>• 8.4</li><li>• 8.5</li><li>• 8.6 &lt;- default</li><li>• 8.7</li></ul>  |
| <b>iso_file_path</b><br>string<br>Required        | Path where user has placed the iso image that needs to be provisioned on target nodes. Accepted files are Rocky8-DVD or RHEL-8.x-DVD (full OS). iso_file_path should contain the provision_os and provision_os_version values in the filename.<br><b>Default values:</b> "/home/RHEL-8.6.0-20220420.3-x86_64-dvd1.iso" |
| <b>timezone</b><br>string<br>Required             | Timezone to be used during OS provisioning. Available timezones are provided <a href="#">here</a> .<br>Choices: <ul style="list-style-type: none"><li>• GMT &lt;- default</li><li>• EST</li><li>• CET</li><li>• MST</li><li>• CST6CDT</li><li>• PST8PDT</li></ul>  |
| <b>language</b><br>string<br>Required             | Language to be used during OS provisioning.<br><b>Default values:</b> en-US  |
| <b>default_lease_time</b><br>integer<br>Required  | Default lease time for IPs assigned by DHCP. Range: 21600-86400<br><b>Default values:</b> 86400  |
| <b>provision_password</b><br>string<br>Required   | <ul style="list-style-type: none"><li>• Password set for the root account of target nodes during provisioning.</li><li>• Length &gt;= 8 characters</li><li>• Password must not contain -, , ‘, ’</li><li>• The first character of the string should be an alphabet.</li></ul>  |
| <b>postgresdb_password</b><br>string<br>Required  | <ul style="list-style-type: none"><li>• Password set for the postgresDB on target nodes during provisioning.</li><li>• Length &gt;= 8 characters</li><li>• Password must not contain -, , ‘, ’</li><li>• The first character of the string should be an alphabet.</li></ul>  |

continues on next page



Table 3 – continued from previous page

| Parameter                                     | Details  |
|---|--|
| <b>node_name</b><br>string<br>Required        | <ul style="list-style-type: none"> <li>Prefix for target node names, if dynamically allocated.</li> <li>Hostname = node_name + '0000x' + domain_name</li> <li>Hostname &lt;= 65 characters</li> <li>Example: servernode00001.Omnia.test , where node_name =servernode, domain_name =Omnia.test , 00001 used by Omnia.</li> </ul> <b>Default values:</b> node   |
| <b>domain_name</b><br>string<br>Required      | <ul style="list-style-type: none"> <li>Domain name the user intends to configure on the cluster.</li> <li>Hostname = node_name + '0000x' + domain_name</li> <li>Hostname &lt;= 65 characters</li> <li>Please provide a valid domain name according to the domain name standards.</li> <li>Example: servernode00001.Omnia.test , where node_name=servernode, domain_name=Omnia.test , 00001 used by Omnia.</li> </ul> |
| <b>public_nic</b><br>string<br>Required       | The nic/ethernet card that is connected to the public internet.<br><b>Default values:</b> eno2   |
| <b>admin_nic</b><br>string<br>Required        | Admin NIC of Control Plane. This is the shared LOM NIC.<br><b>Default values:</b> eno1   |
| <b>admin_nic_subnet</b><br>string<br>Required | <ul style="list-style-type: none"> <li>The subnet within which all Admin IPs are assigned.</li> <li>The value of this variable cannot be changed after successfully running provision.yml.</li> </ul> <b>Default values:</b> 10.5.0.0  |
| <b>ib_nic_subnet</b><br>string<br>Optional    | <ul style="list-style-type: none"> <li>If provided, Omnia will handle and assign static IPs to cluster node's IB network.</li> <li>Only the last 16 bits/2 octets of IPv4 are dynamic</li> <li>If provided, the DB entry will be in parallel with the pxe_subnet.</li> <li>Example: If admin_ip: 10.5.0.50 and ib_nic_subnet: 10.10.0.0, then ib_ip: 10.10.0.50</li> </ul>   |

continues on next page

Table 3 – continued from previous page

| Parameter   | Details  |
|---|--|
| <b>bmc_nic_subnet</b><br>string<br>Required         | <ul style="list-style-type: none"> <li>• If provided, Omnia will assign static IPs to BMC NICs on the cluster nodes within the provided subnet.</li> <li>• If <code>network_interface_type</code>: <code>lom</code>, mandatory for <code>discovery_mechanism</code>: <code>mapping</code>, <code>switch_based</code> and <code>bmc</code>.</li> <li>• If <code>network_interface_type</code>: <code>dedicated</code>, optional for <code>discovery_mechanism</code>: <code>mapping</code>, <code>snmp-walk</code>.</li> <li>• Note that since the last 16 bits/2 octets of IPv4 are dynamic, please ensure that the parameter value is set to <code>x.x.0.0</code>.</li> <li>• When the PXE range and BMC subnet are provided, corresponding NICs will be assigned IPs with the same 3rd and 4th octets.</li> <li>• The value of this variable cannot be changed after successfully running <code>provision.yml</code>.</li> </ul>   |
| <b>bmc_username</b><br>string<br>Required           | <ul style="list-style-type: none"> <li>• The username for iDRAC.</li> <li>• The username must not contain <code>-</code>, <code>,</code></li> </ul>  |
| <b>bmc_password</b><br>string<br>Required           | <ul style="list-style-type: none"> <li>• The password for iDRAC.</li> <li>• The password must not contain <code>-</code>, <code>,</code></li> <li>• The first character of the string should be an alphabet.</li> </ul>  |
| <b>bmc_static_start_range</b><br>string<br>Required | <ul style="list-style-type: none"> <li>• The dhcp range for discovering the static IP assigned iDRACs within the given range.</li> <li>• For the range <code>10.3.0.50</code> to <code>10.3.4.100</code> then, <code>bmc_static_start_range</code>: <code>10.3.0.50</code>, <code>bmc_static_end_range</code>: <code>10.3.4.100</code></li> <li>• To create a meaningful range of discovery, ensure that the last two octets of <code>bmc_static_end_range</code> are equal to or greater than the last two octets of the <code>bmc_static_start_range</code>. That is, for the range <code>a.b.c.d - a.b.e.f</code>, <code>e</code> and <code>f</code> should be greater than or equal to <code>c</code> and <code>d</code>.</li> <li>• The value of this variable cannot be changed after successfully running <code>provision.yml</code>.</li> <li>• Ex: <code>172.20.0.50 - 172.20.1.101</code> is a valid range however, <code>172.20.0.101 - 172.20.1.50</code> is not.</li> </ul> |

continues on next page

Table 3 – continued from previous page

| Parameter  | Details  |
|--|--|
| <b>bmc_static_end_range</b><br>string<br>Required                                  | <ul style="list-style-type: none"> <li>• The dhcp range for discovering the static IP assigned iDRACs within the given range.</li> <li>• For the range 10.3.0.50 to 10.3.4.100 then, bmc_static_start_range: 10.3.0.50, bmc_static_end_range: 10.3.4.100</li> <li>• To create a meaningful range of discovery, ensure that the last two octets of bmc_static_end_range are equal to or greater than the last two octets of the bmc_static_start_range. That is, for the range a.b.c.d - a.b.e.f, e and f should be greater than or equal to c and d.</li> <li>• The value of this variable cannot be changed after successfully running provision.yml.</li> <li>• Ex: 172.20.0.50 - 172.20.1.101 is a valid range however, 172.20.0.101 - 172.20.1.50 is not.</li> </ul> |
| <b>update_repos</b><br>boolean <sup>1</sup><br>Required                            | <ul style="list-style-type: none"> <li>• Indicates whether provision.yml will update offline RHEL repos.</li> <li>• If update_repos: false, the update repos for BaseOS and AppStream will not be updated to the latest versions available.</li> <li>• If update_repos: true, the update repos for BaseOS and AppStream will be updated to the latest versions available.</li> </ul> <hr/> <p><b>Note:</b> By default, AppStream and BaseOS repos will be configured from the given ISO file.</p> <hr/> <p>Choices:<br/> false &lt;- Default<br/> true</p>   |
| <b>rhel_repo_alphabetical_folders</b><br>boolean <sup>Page 45, 1</sup><br>Required | <ul style="list-style-type: none"> <li>• Indicates whether the packages in local or subscription repos should be ordered in alphabetical directories.</li> <li>• This variable should be filled if control plane OS is RHEL and local RHEL repository is available.</li> </ul> <p>Choices:<br/> false &lt;- Default<br/> true</p>  |

continues on next page

Table 3 – continued from previous page

| Parameter  | Details   |
|--|---|
| <b>rhel_repo_local_path</b><br>JSON list<br>Optional | <ul style="list-style-type: none"> <li>• The repo path and names of the software repository to be configured on the cluster nodes.</li> <li>• Provide the repo data file path, which ends with .repo extension in repo_url parameter.</li> <li>• Provide a <b>valid</b> url for BaseOS and AppStream repositories. Omnia does not validate the repo_url provided. Invalid entries will cause provision.yml to fail.</li> <li>• This variable should be filled if control plane OS is RHEL and subscription is not activated.</li> <li>• This variable should be filled if the control plane OS is Rocky and the provision_os is rhel.</li> </ul> <p>Sample value:</p> <pre>- { repo: "AppStream", repo_url: "http://xx.yy.zz/pub/ Distros/RedHat/RHEL8/8.8/RHEL-8-appstream.repo", repo_name: "RHEL-8-appstream-partners" } - { repo: "BaseOS", repo_url: "http://xx.yy.zz/pub/ Distros/RedHat/RHEL8/8.8/RHEL-8-baseos.repo", repo_name: "RHEL-8-baseos-partners" }</pre> |
| <b>primary_dns</b><br>string<br>Optional             | <ul style="list-style-type: none"> <li>• The primary DNS host IP queried to provide Internet access to cluster Node (through DHCP routing).</li> <li>• Currently, the primary_dns value stored in input/provision_config.yml cannot be part of any of the subnets (admin_nic_subnet, ib_nic_subnet and bmc_nic_subnet) also defined in input/provision_config.yml.</li> </ul> <p>Ex: If the primary_dns is set to 10.15.0.7, the subnet 10.15.0.0 cannot be used for admin_nic_subnet, ib_nic_subnet or bmc_nic_subnet.</p>   |
| <b>secondary_dns</b><br>string<br>Optional           | <p>The secondary DNS host IP queried to provide Internet access to cluster Node (through DHCP routing)</p>  |

continues on next page

Table 3 – continued from previous page

| Parameter   | Details   |
|---|---|
| <b>disk_partition</b><br>JSON list<br>Optional                        | <ul style="list-style-type: none"> <li>• User defined disk partition applied to remote servers.</li> <li>• The disk partition desired_capacity has to be provided in MB.</li> <li>• Valid mount_point values accepted for disk partition are /home, /var, /tmp, /usr, swap.</li> <li>• Default partition size provided for /boot is 1024MB, /boot/efi is 256MB and the remaining space to / partition.</li> <li>• Values are accepted in the form of JSON list such as: , - { mount_point: “/home”, desired_capacity: “102400” }</li> </ul> <p><b>Default values:</b> - { mount_point: “”, desired_capacity: “” }</p> |
| <b>mlnx_ofed_path</b><br>string<br>Optional                           | Absolute path to a local copy of the .iso file containing Mellanox OFED packages. The image can be downloaded from <a href="https://network.nvidia.com/products/infiniband-drivers/linux/mlnx_ofed/">https://network.nvidia.com/products/infiniband-drivers/linux/mlnx_ofed/</a> . Sample value: /root/MLNX_OFED_LINUX-5.8-1.1.2.1-rhel8.6-x86_64.iso   |
| <b>cuda_toolkit_path</b><br>string<br>Optional                        | Absolute path to local copy of .rpm file containing CUDA packages. The cuda rpm can be downloaded from <a href="https://developer.nvidia.com/cuda-downloads">https://developer.nvidia.com/cuda-downloads</a> . CUDA will be installed post provisioning without any user intervention. Eg: cuda_toolkit_path: “/root/cuda-repo-rhel8-12-0-local-12.0.0-525.60.13-1.x86_64.rpm”  |
| <b>apptainer_support</b><br>boolean <sup>Page 45, 1</sup><br>Required | <ul style="list-style-type: none"> <li>• Indicates whether apptainer will be installed on the cluster to enable execution of HPC benchmarks in a containerized environment.</li> <li>• If apptainer_support: false, apptainer will not be installed on the cluster.</li> <li>• If apptainer_support: true, apptainer will be installed on the cluster.</li> </ul>   |

**Note:** The input/provision\_config.yml file is encrypted on the first run of the provision tool:

To view the encrypted parameters:

```
ansible-vault view provision_config.yml --vault-password-file .provision_vault_
↪key
```

To edit the encrypted parameters:

```
ansible-vault edit provision_config.yml --vault-password-file .provision_vault_
↪key
```

<sup>1</sup> Boolean parameters do not need to be passed with double or single quotes.

**Caution:** The IP address *192.168.25.x* is used for PowerVault Storage communications. Therefore, do not use this IP address for other configurations.

To continue to the next steps:

- [Provisioning the cluster](#)

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

## snmpwalk

Omnia can query known switches (by IP and community string) for information on target node MAC IDs. The following parameters need to be populated in `input/provision_config.yml` to discover target nodes using SNMP.

### Pre requisites

SNMPv2 should be enabled on the switch specified using `pxe_switch_ip` in `input/provision_config`.

To enable SNMPv2, log in to the switch and run the following commands:

```
configure terminal
snmp-server community public ro
exit
```

Use `show snmp community` to verify your changes.

**Note:** The commands provided above sets the SNMP community string of the switch to `public`. Ensure that the community string set above matches the value provided in `pxe_switch_snmp_community_string` in `input/provision_config.yml`

### Caution:

- Target servers with LOM architecture is not supported.
- Do not remove or comment any lines in the `input/provision_config.yml` file.
- `admin_nic_subnet`, `ib_nic_subnet` and `bmc_nic_subnet` should have the same subnet mask (Omnia only supports /16 subnet masks currently).
- **THE ROCKY LINUX OS VERSION ON THE CLUSTER WILL BE UPGRADED TO THE LATEST 8.x VERSION AVAILABLE IRRESPECTIVE OF THE PROVISION\_OS\_VERSION PROVIDED IN PROVISION\_CONFIG.YML.**

Table 4: Parameters

| Parameter   | Details  |
|---|--|
| <b>network_interface_type</b><br>string<br>Required | The network type used on the Omnia cluster.<br>Choices: <ul style="list-style-type: none"><li>• dedicated &lt;- default</li><li>• lom</li></ul>  |
| <b>discovery_mechanism</b><br>string<br>Required    | The mechanism through which Omnia will discover nodes for provisioning. For more information on how the mechanisms work, go to <a href="#">DiscoveryMechanisms</a> .<br>Choices: <ul style="list-style-type: none"><li>• switch_based</li><li>• mapping &lt;-default</li><li>• bmc</li><li>• snmpwalk</li></ul>                                |
| <b>provision_os</b><br>string<br>Required           | The operating system to be provisioned on target nodes in the cluster.<br>Choices: <ul style="list-style-type: none"><li>• rhel &lt;-default</li><li>• rocky</li></ul> <div><b>Caution:</b><br/>THE ROCKY LINUX OS VERSION ON THE CLUSTER WILL BE UP-GRADED TO THE LATEST 8.x VERSION AVAILABLE IRRESPECTIVE OF THE PROVISION_OS_VERSION</div> |

**Caution:** The IP address *192.168.25.x* is used for PowerVault Storage communications. Therefore, do not use this IP address for other configurations.

---

**Note:** The input/provision\_config.yml file is encrypted on the first run of the provision tool:

To view the encrypted parameters:

```
ansible-vault view provision_config.yml --vault-password-file .provision_vault_  
↪key
```

To edit the encrypted parameters:

```
ansible-vault edit provision_config.yml --vault-password-file .provision_vault_  
↪key
```

---

To continue to the next steps:

- [Provisioning the cluster](#)

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

### switch\_based

Omnia can query known switches (by SNMPv3 username/password) for information on target node MAC IDs.

#### Pros

- The whole discovery process is totally automatic.
- Admin IP, BMC IP and Infiniband IP address configuration is automatic on the target nodes.
- Re-provisioning of servers will be automatic.
- PXE booting servers is supported via split ports on the switch.

#### Cons

- Users need to enable IPMI on target servers.
- Servers require a manual PXE boot after the first run of the provision tool

For more information regarding switch-based discovery, [click here](#)

### mapping

Manually collect PXE NIC information for target servers and manually define them to Omnia using a mapping file using the below format:

#### pxe\_mapping\_file.csv

```
MAC,Hostname,IP  
  
xx:yy:zz:aa:bb:cc,server,10.5.0.101  
aa:bb:cc:dd:ee:ff,server2, 10.5.0.102
```

#### Pros

- Easily customized if the user maintains a list of MAC addresses.

---

<sup>1</sup> Boolean parameters do not need to be passed with double or single quotes.



**Cons**

- The user needs to be aware of the MAC/IP mapping required in the network.
- Servers require a manual PXE boot if iDRAC IPs are not configured.

For more information regarding mapping files, [click here](#)

**bmc**

Omnia can also discover nodes via their iDRAC using IPMI.

**Pros**

- Discovery and provisioning of servers is automatic.
- Admin, BMC and Infiniband IP address configuration is automatic on the control plane.
- LOM architecture is supported (including cloud enclosures: C6420, C6520, C6620).

**Cons**

- For iDRACs that are not DHCP enabled (ie Static), users need to enable IPMI manually.

For more information regarding BMC, [click here](#)

**snmpwalk**

Omnia can query known switches (by IP and community string) for information on target node MAC IDs.

**Pros**

- The method can be applied to large clusters.
- User intervention is minimal.

**Cons**

- Switches should be SNMP enabled.
- Servers require a manual PXE boot if iDRAC IPs are not configured.
- PXE NIC ranges should contain IPs that are double the iDRACs present (as NIC and iDRAC MACs may need to be mapped).
- LOM architecture is not supported (including cloud enclosures: C6420, C6520, C6620).

For more information regarding snmpwalk, [click here](#)

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

## 2.2.3 Provisioning the cluster

Edit the `input/provision_config.yml` file to update the required variables. A list of the variables required is available by [discovery mechanism](#).

---

**Note:** The first PXE device on target nodes should be the designated active NIC for PXE booting.

---

### Network Settings

|   | Current Value                         |
|---|---------------------------------------|
| PXE Device1   | Enabled ▾                             |
| PXE Device2   | Disabled ▾                            |
| PXE Device3   | Disabled ▾                            |
| PXE Device4   | Disabled ▾                            |
| <h4> <div> <div></div> <div>PXE Device1 Settings</div> </div> </h4> |                                       |
|   | Current Value                         |
| Interface   | Integrated NIC 1 Port 1 Partition 1 ▾ |
| Protocol  | IPv4 ▾                                |
| VLAN  | Enabled ▾                             |

## Optional configurations managed by the provision tool

### Installing CUDA

#### Using the provision tool

- If `cuda_toolkit_path` is provided in `input/provision_config.yml` and NVIDIA GPUs are available on the target nodes, CUDA packages will be deployed post provisioning without user intervention.

#### Using the Accelerator playbook

- CUDA can also be installed using `accelerator.yml` after provisioning the servers (Assuming the provision tool did not install CUDA packages).

---

#### Note:

- The CUDA package can be downloaded from [here](#)
  - CUDA requires an additional reboot while being installed. While this is taken care of by Omnia, users are required to wait an additional few minutes when running the provision tool with CUDA installation for the target nodes to come up.
- 

### Installing OFED

#### Using the provision tool

- If `mlnx_ofed_path` is provided in `input/provision_config.yml` and Mellanox NICs are available on the target nodes, OFED packages will be deployed post provisioning without user intervention.

---

**Note:** When leveraging the provision tool to install OFED, Omnia prevents the following packages from being upgraded: `dapl*` `ibacm` `infiniband*` `libmlx*` `librdma*` `opensm*` `ibutils*` `perftest*` `openmpi` by appending `exclude=dapl* ibacm infiniband* libmlx* librdma* opensm* ibutils* perftest* openmpi` to `/etc/yum.conf`. For more information on this, [click here](#).

---

### Using the Network playbook

- OFED can also be installed using [network.yml](#) after provisioning the servers (Assuming the provision tool did not install OFED packages).

---

**Note:** The OFED package can be downloaded from [here](#) .

---

### Assigning infiniband IPs

When `ib_nic_subnet` is provided in `input/provision_config.yml`, the infiniband NIC on target nodes are assigned IPv4 addresses within the subnet without user intervention. When PXE range and Infiniband subnet are provided, the infiniband NICs will be assigned IPs with the same 3rd and 4th octets as the PXE NIC.

- For example on a target node, when the PXE NIC is assigned 10.5.0.101, and the Infiniband NIC is assigned 10.10.0.101 (where `ib_nic_subnet` is 10.10.0.0).

---

**Note:** The IP is assigned to the interface **ib0** on target nodes only if the interface is present in **active** mode. If no such NIC interface is found, xCAT will list the status of the node object as failed.

---

### Assigning BMC IPs

When target nodes are discovered via SNMP or mapping files (ie `discovery_mechanism` is set to `snmp` or `mapping` in `input/provision_config.yml`), the `bmc_nic_subnet` in `input/provision_config.yml` can be used to assign BMC IPs to iDRAC without user intervention. When PXE range and BMC subnet are provided, the iDRAC NICs will be assigned IPs with the same 3rd and 4th octets as the PXE NIC.

- For example on a target node, when the PXE NIC is assigned 10.5.0.101, and the iDRAC NIC is assigned 10.3.0.101 (where `bmc_nic_subnet` is 10.3.0.0).

### Using multiple versions of a given OS

Omnia now supports deploying different versions of the same OS. With each run of `provision.yml`, a new deployable OS image is created with a distinct type (rocky or RHEL) and version (8.0, 8.1, 8.2, 8.3, 8.4, 8.5, 8.6, 8.7) depending on the values provided in `input/provision_config.yml`.

---

#### Note:

- While Omnia deploys the minimal version of the OS, the multiple version feature requires that the Rocky full (DVD) version of the OS be provided.
  - The multiple OS feature is only available with Rocky 8.7 when xCAT 2.16.5 is in use. [Currently, Omnia uses 2.16.4]
- 

### DHCP routing for internet access

Omnia now supports DHCP routing via the control plane. To enable routing, update the `primary_dns` and `secondary_dns` in `input/provision_config.yml` with the appropriate IPs (hostnames are currently not supported). For cluster nodes that are not directly connected to the internet (ie only PXE network is configured), this configuration allows for internet connectivity.

### Disk partitioning

Omnia now allows for customization of disk partitions applied to remote servers. The disk partition `desired_capacity` has to be provided in MB. Valid `mount_point` values accepted for disk partition

are /home, /var, /tmp, /usr, swap. Default partition size provided for /boot is 1024MB, /boot/efi is 256MB and the remaining space to / partition. Values are accepted in the form of JSON list such as:

```
disk_partition:
- { mount_point: "/home", desired_capacity: "102400" }
- { mount_point: "swap", desired_capacity: "10240" }
```

## Running the provision tool

To deploy the Omnia provision tool, run the following command

```
cd provision
ansible-playbook provision.yml
```

provision.yml runs in three stages that can be called individually:

### Preparing the control plane

- Verifies pre-requisites such as SELinux and xCAT services status.
- Installs required tool packages.
- Verifies and updates firewall settings.
- Installs xCAT.
- Configures xCAT databases basis input/provision\_config.yml.
- Configures a docker registry to pull images from the internet and store them locally. These images will then be pulled by cluster nodes locally. For more information, [click here](#).

To call this playbook individually, ensure that input/provision\_config.yml is updated and then run:

```
ansible-playbook prepare_cp.yml
```

### Creating/updating the repositories

- Creates and updates all repositories required locally.
- This playbook also invokes the airgap.yml playbook for creating offline repositories. For more information on this, [click here](#).
- To call this playbook individually, ensure that prepare\_cp.yml has run at least once and then run:

```
ansible-playbook repo_manipulate.yml
```

### Discovering/provisioning the nodes

- Discovers all target servers based on specifications in input/provision\_config.yml.
- PostgreSQL database is set up with all relevant cluster information such as MAC IDs, hostname, admin IP, infiniband IPs, BMC IPs etc.
- Configures the control plane with NTP services for cluster node synchronization.
- Provisions all discovered servers.

To call this playbook individually, ensure that repo\_manipulate.yml has run at least once and then run:

```
ansible-playbook discovery_provision.yml
```

**After successfully running `provision.yml`, go to Building Clusters to setup Slurm, Kubernetes, NFS, BeeGFS and Authentication.**

**Note:**

- Ansible playbooks by default run concurrently on 5 nodes. To change this, update the `forks` value in `ansible.cfg` present in the respective playbook directory.
- While the `admin_nic` on cluster nodes is configured by Omnia to be static, the public NIC IP address should be configured by user.
- If the target nodes were discovered using switch-based, mapping or `snmpwalk` mechanisms, manually PXE boot the target servers after the `provision.yml` playbook is executed and the target node lists as **booted** in the `nodeinfo` table.
- If the cluster does not have access to the internet, AppStream will not function. To provide internet access through the control plane (via the PXE network NIC), update `primary_dns` and `secondary_dns` in `provision_config.yml` and run `provision.yml`
- All ports required for xCAT to run will be opened (For a complete list, check out the [Security Configuration Document](#)).
- After running `provision.yml`, the file `input/provision_config.yml` will be encrypted. To edit the file, use the command: `ansible-vault edit provision_config.yml --vault-password-file .provision_vault_key`
- To re-provision target servers `provision.yml` can be re-run with a new inventory file that contains a list of admin (PXE) IPs. For more information, [click here](#)
- Post execution of `provision.yml`, IPs/hostnames cannot be re-assigned by changing the mapping file. However, the addition of new nodes is supported as explained [here](#).

**Caution:**

- Once xCAT is installed, restart your SSH session to the control plane to ensure that the newly set up environment variables come into effect.
- To avoid breaking the passwordless SSH channel on the control plane, do not run `ssh-keygen` commands post execution of `provision.yml` to create a new key.
- **Do not delete the following directories:**

- `/root/xcat`
- `/root/xcat-dbbback`
- `/docker-registry`
- `/opt/omnia`
- `/var/log/omnia`

- On subsequent runs of `provision.yml`, if users are unable to log into the server, refresh the ssh key manually and retry.

```
ssh-keygen -R <node IP>
```

- If a subsequent run of `provision.yml` fails, the `input/provision_config.yml` file will be unencrypted.

To create a node inventory in /opt/omnia, [click here](#).

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

## 2.2.4 Checking node status

### Via CLI

Run `nodels all nodelist.status` for a list of nodes and their statuses.

```
omnia-node00001: installing
omnia-node00002: booted
omnia-node00003: powering-on
omnia-node00004: booted
```

Possible values of node status are powering-off, powering-on, bmcready, installing, booting, post-booting, booted, failed.

**Caution:** Once xCAT is installed, restart your SSH session to the control plane to ensure that the newly set up environment variables come into effect. This will also allow the above command to work correctly.

### Via omniadb

1. To access the DB, run:

```
psql -U postgres
\c omniadb
```

2. To view the schema being used in the cluster: `\dn`
3. To view the tables in the database: `\dt`
4. To view the contents of the nodeinfo table: `select * from cluster.nodeinfo;`

| id | serial            | node            | hostname                   | admin_mac  |
|----|-------------------|-----------------|----------------------------|------------|
|    | admin_ip          | bmc_ip          | ib_ip                      | status     |
|    | switch_ip         | switch_name     | switch_port                | bmc_mode   |
| 1  | XXXXXXX           | control_plane   | control_plane.omnia.test   |            |
|    | ec:2a:72:34:f7:26 | 10.5.255.254    | 10.19.255.254              |            |
| 2  | XXXXXXX           | omnia-node00002 | omnia-node00001.omnia.test |            |
|    | 10.5.0.102        | 10.19.0.102     | 10.10.0.102                | booted     |
|    | 28.132            | switch1         | 3                          |            |
| 3  | XXXXXXX           | omnia-node00003 | omnia-node00002.omnia.test |            |
|    | 10.5.0.103        | 10.19.0.103     | 10.10.0.103                |            |
|    | 28.132            | switch1         | 4                          |            |
| 4  | XXXXXXX           | omnia-node00004 | omnia-node00003.omnia.test |            |
|    | 2c:ea:7f:3d:6b:98 | 10.5.0.104      | 10.19.0.104                | installing |

(continues on next page)



## 2.2.5 Configuring servers with out-of-band management

For pre-configured iDRACs, `provision/idrac.yml` can be used to provision the servers.

### Before running `idrac.yml`

- The `idrac_inventory` file is updated with the iDRAC IP addresses.
- To customize iDRAC provisioning, input parameters can be updated in the `provision/idrac_input.yml` file.
- The Lifecycle Controller Remote Services of PowerEdge Servers is in the 'ready' state.
- The Redfish services are enabled in the iDRAC settings under Services.
- The provision tool has discovered the servers using SNMP/mapping.
- iDRAC 9 based Dell EMC PowerEdge Servers with firmware versions 5.00.10.20 and above. (With the latest BIOS available)

### Configurations performed by `idrac.yml`

- If bare metal servers have BOSS controllers installed, virtual disks (Data will be stored in a RAID 1 configuration by default) will be created on the BOSS controller (ie, RAID controllers will be ignored/unmanaged). Ensure that exactly 2 SSD disks are available on the server.
- If bare metal servers have a RAID controller installed, Virtual disks are created for RAID configuration (Data will be saved in a RAID 0 configuration by default).
- Omnia validates and configures the active host NICs in PXE device settings when `provision_method` is set to PXE. (If no active NIC is found, `idrac.yml` will fail on the target node.)
- Once all configurations are in place, the `idrac.yml` initiates a PXE boot for configuration to take effect.

---

### Note:

- Servers that have not been discovered by the Provision tool will not be provisioned with the OS image.
  - Since the BMC discovery method PXE boots target iDRACs while running the provision tool, this script is not recommended for such servers.
- 

### Running `idrac.yml`

```
ansible-playbook idrac.yml -i idrac_inventory -e idrac_username='' -e idrac_password=''
```

Where the `idrac_inventory` points to the file mentioned above and the `idrac_username` and `idrac_password` are the credentials used to authenticate into iDRAC.

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).



## 2.3 Creating node inventory

Once the **servers are provisioned**, a list of reachable nodes can be fetched using the below command:

```
ansible-playbook post_provision.yml
```

This creates a node inventory in `/opt/omnia`.

```
cat /opt/omnia/node_inventory
10.5.0.100 service_tag=XXXXXXX operating_system=RedHat
10.5.0.101 service_tag=XXXXXXX operating_system=RedHat
10.5.0.102 service_tag=XXXXXXX operating_system=Rocky
10.5.0.103 service_tag=XXXXXXX operating_system=Rocky
```

To create an inventory when **Building Clusters**, use the reachable nodes' IP addresses from the above output to assign manager, compute and login groups. For more information on the inventory file used, [click here](#).

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

## 2.4 Configuring the cluster

### Features enabled by omnia.yml

- Centralized authentication: Once all the required parameters in `security_config.yml` are filled in, `omnia.yml` can be used to set up FreeIPA/LDAP.
- Slurm: Once all the required parameters in `omnia_config.yml` are filled in, `omnia.yml` can be used to set up slurm.
- Login Node (Additionally secure login node)
- Kubernetes: Once all the required parameters in `omnia_config.yml` are filled in, `omnia.yml` can be used to set up kubernetes.
- BeeGFS bolt on installation: Once all the required parameters in `storage_config.yml` are filled in, `omnia.yml` can be used to set up NFS.
- NFS bolt on support: Once all the required parameters in `storage_config.yml` are filled in, `omnia.yml` can be used to set up BeeGFS.
- Telemetry: Once all the required parameters in `telemetry_config.yml` are filled in, `omnia.yml` sets up **Omnia telemetry** and/or **iDRAC telemetry**. It also installs **Grafana** and **Loki** as Kubernetes pods.

### 2.4.1 Input parameters for the cluster

These parameters are located in `input/omnia_config.yml`, `input/security_config.yml`, `input/telemetry_config.yml` and [optional] `input/storage_config.yml`.

**Caution:** Do not remove or comment any lines in the `input/omnia_config.yml`, `input/security_config.yml` and [optional] `input/storage_config.yml` file.

`omnia_config.yml`

Table 5: Parameters for kubernetes

| Variables   | Details   |
|---|---|
| <b>scheduler_type</b><br>string<br>Required           | <ul style="list-style-type: none"> <li>Job scheduler to be installed across all nodes in the cluster*</li> <li>To install slurm provide scheduler_type: "slurm"</li> <li>To install k8s provide scheduler_type: "k8s"</li> <li>To install slurm and k8s provide scheduler_type: "slurm,k8s"</li> </ul> <b>Default value:</b> "slurm"  |
| <b>k8s_version</b><br>string<br>Required              | <ul style="list-style-type: none"> <li>Kubernetes version.</li> <li><b>Required when scheduler_type: "k8s"</b><br/> Choices: <ul style="list-style-type: none"> <li>"1.19.3" &lt;- default</li> <li>" 1.16.7"</li> </ul> </li> </ul>  |
| <b>k8s_cni</b><br>string<br>Required                  | <ul style="list-style-type: none"> <li>Kubernetes SDN network.</li> <li>Required when scheduler_type: "k8s"<br/> Choices: <ul style="list-style-type: none"> <li>"calico" &lt;- default</li> <li>"flannel"</li> </ul> </li> </ul>   |
| <b>k8s_pod_network_cidr</b><br>string<br>Required     | <ul style="list-style-type: none"> <li>Kubernetes pod network CIDR.</li> <li>Make sure this value does not overlap with any of the host networks.</li> <li>Required when scheduler_type: "k8s"<br/> <b>Default values:</b> "10.244.0.0/16"</li> </ul>   |
| <b>docker_username</b><br>string<br>Optional          | <ul style="list-style-type: none"> <li>Username for Dockerhub account</li> <li>A kubernetes secret will be created and patched to service account in default namespace. This kubernetes secret can be used to pull images from private repositories</li> <li>This value is optional but suggested avoiding docker pull limit issues.</li> <li>The first character of the string should be an alphabet.</li> </ul> |
| <b>docker_password</b><br>string<br>Optional          | <ul style="list-style-type: none"> <li>Password for Dockerhub account</li> <li>This value is mandatory if docker username is provided.</li> <li>The first character of the string should be an alphabet.</li> </ul>   |
| <b>ansible_config_file_path</b><br>string<br>Required | <ul style="list-style-type: none"> <li>Path to directory hosting ansible config file (ansible.cfg file)</li> <li>This directory is on the host running ansible, if ansible is installed using dnf</li> <li>If ansible is installed using pip, this path should be set</li> </ul>  |
| <b>enable_omnia_nfs</b><br>boolean <sup>1</sup>       | <ul style="list-style-type: none"> <li>Boolean indicating whether a parallel file system</li> </ul>   |

Table 6: Parameters for slurm setup

| Variables  | Details   |
|--|---|
| <b>scheduler_type</b><br>string<br>Required                          | <ul style="list-style-type: none"> <li>• Job scheduler to be installed across all nodes in the cluster*</li> <li>• To install slurm provide scheduler_type: "slurm"</li> <li>• To install k8s provide scheduler_type: "k8s"</li> <li>• To install slurm and k8s provide scheduler_type: "slurm,k8s"</li> </ul> <b>Default value:</b> "slurm"  |
| <b>mariadb_password</b><br>string<br>Optional                        | <ul style="list-style-type: none"> <li>• Password used for Slurm database.</li> <li>• The password must not contain -, ,</li> <li>• The length of the password should be at least 8.</li> <li>• Required when scheduler_type: "slurm".</li> <li>• The first character of the string should be an alphabet.</li> </ul> <b>Default value:</b> "password"  |
| <b>ansible_config_file_path</b><br>string<br>Required                | <ul style="list-style-type: none"> <li>• Path to directory hosting ansible config file (ansible.cfg file)</li> <li>• This directory is on the host running ansible, if ansible is installed using dnf</li> <li>• If ansible is installed using pip, this path should be set</li> </ul> <b>Default values:</b> /etc/ansible  |
| <b>enable_omnia_nfs</b><br>boolean <sup>Page 59, 1</sup><br>Required | <ul style="list-style-type: none"> <li>• Boolean indicating whether a parallel file system is <b>not</b> running in the environment and a share file system (NFS/BeeGFS) will be used to create home directory/ Kubernetes share directory on it.</li> <li>• When this variable is <b>true</b>, Omnia will create its own NFS share and mount omnia_usrhome_share on all the nodes.</li> </ul> Choices: <ul style="list-style-type: none"> <li>– true &lt;- default</li> <li>– false</li> </ul> |
| <b>omnia_usrhome_share</b><br>string<br>Required                     | <ul style="list-style-type: none"> <li>• Path to directory which will be shared across all nodes in the cluster.</li> <li>• If enable_omnia_nfs: true, NFS share will be created at path mentioned below.</li> <li>• If enable_omnia_nfs: false, set this variable as path to parallel file system(NFS/BeeGFS) running in the system.</li> </ul> <b>Default value:</b> "/home/omnia-share"  |

security\_config.yml

<sup>1</sup> Boolean parameters do not need to be passed with double or single quotes.

Table 7: Parameters for FreeIPA

| Parameter  | Details   |
|--|---|
| <b>freeipa_required</b><br>boolean <sup>1</sup> Required | Boolean indicating whether FreeIPA is required or not.<br>Choices:<br><ul style="list-style-type: none"> <li>• <code>true</code> &lt;- Default</li> <li>• <code>false</code></li> </ul>   |
| <b>realm_name</b><br>string Required                     | Sets the intended realm name.<br><b>Default value:</b> <code>OMNIA.TEST</code>  |
| <b>directory_manager_password</b><br>string Required     | <ul style="list-style-type: none"> <li>• Password authenticating admin level access to the Directory for system management tasks.</li> <li>• It will be added to the instance of directory server created for IPA.</li> <li>• Required Length: 8 characters.</li> <li>• The password must not contain <code>-</code>, <code>,</code> ,</li> <li>• The first character of the string should be an alphabet.</li> </ul> |
| <b>kerberos_admin_password</b><br>string Required        | <ul style="list-style-type: none"> <li>• “admin” user password for the IPA server on RockyOS.</li> <li>• The first character of the string should be an alphabet.</li> </ul>  |
| <b>domain_name</b><br>string Required                    | Sets the intended domain name<br><b>Default value:</b> <code>omnia.test</code>  |

Table 8: Parameters for LDAP

| Parameter   | Details   |
|---|---|
| <b>ldap_required</b><br>boolean <small>Page 59, 1</small> Optional            | Boolean indicating whether ldap client is required or not.<br>Choices:<br><ul style="list-style-type: none"> <li>• false &lt;- Default</li> <li>• true</li> </ul>   |
| <b>domain_name</b><br>string Optional   | Sets the intended domain name.<br><b>Default values:</b> omnia.test   |
| <b>ldap_server_ip</b><br>string Optional                                      | LDAP server IP. Required if ldap_required is true. There should be an explicit LDAP server running on this IP.  |
| <b>ldap_connection_type</b><br>string Optional                                | For a TLS connection, provide a valid certification path. For an SSL connection, ensure port 636 is open.<br><b>Default values:</b> TLS   |
| <b>ldap_ca_cert_path</b><br>string Optional                                   | This variable accepts Server Certificate Path. Make sure certificate is present in the path provided. The certificate should have .pem or .crt extension. This variable is mandatory if connection type is TLS.<br><b>Default values:</b> /etc/openldap/certs/omnialdap.pem   |
| <b>user_home_dir</b><br>string Optional                                       | This variable accepts the user home directory path for ldap configuration. If nfs mount is created for user home, make sure you provide the LDAP users mount home directory path.<br><b>Default values:</b> /home   |
| <b>ldap_bind_username</b><br>string Optional                                  | If LDAP server is configured with bind dn then bind dn user to be provided. If this value is not provided (when bind is configured in server) then ldap authentication fails. Omnia does not validate this input. Ensure that it is valid and proper.<br><b>Default values:</b> admin   |
| <b>ldap_bind_password</b><br>string Optional                                  | <ul style="list-style-type: none"> <li>• If LDAP server is configured with bind dn then bind dn password to be provided. If this value is not provided (when bind is configured in server) then ldap authentication fails. Omnia does not validate this input. Ensure that it is valid and proper.</li> <li>• The first character of the string should be an alphabet.</li> </ul> |
| <b>enable_secure_login_node</b><br>boolean <small>Page 59, 1</small> Optional | Boolean value deciding whether security features are enabled on the Login Node.<br>Choices:<br><ul style="list-style-type: none"> <li>• false &lt;- Default</li> <li>• true</li> </ul>  |

**storage\_config.yml**

| Name   | Details   |
|--|---|
| <b>nfs_client_params</b><br>JSON list Optional | <p>If NFS client services are to be deployed, enter the configuration required here in JSON format. The server_ip provided should have an explicit NFS server running. If left blank, no NFS configuration takes place. Possible values include:</p> <ol style="list-style-type: none"> <li>1. Single NFS file system: A single filesystem from a single NFS server is mounted.<br/>           Sample value: <code>- { server_ip: xx.xx.xx.xx, server_share_path: "/mnt/share", client_share_path: "/mnt/client", client_mount_options: "nosuid,rw,sync,hard,intr" }</code></li> <li>2. Multiple Mount NFS file system: Multiple filesystems from a single NFS server are mounted. Sample values:               <ul style="list-style-type: none"> <li>- <code>{ server_ip: xx.xx.xx.xx, server_share_path: "/mnt/server1", client_share_path: "/mnt/client1", client_mount_options: "nosuid,rw,sync,hard,intr" }</code></li> <li>- <code>{ server_ip: xx.xx.xx.xx, server_share_path: "/mnt/server2", client_share_path: "/mnt/client2", client_mount_options: "nosuid,rw,sync,hard,intr" }</code></li> </ul> </li> <li>3. Multiple NFS file systems: Multiple filesystems are mounted from multiple servers. Sample Values:               <ul style="list-style-type: none"> <li>- <code>{ server_ip: zz.zz.zz.zz, server_share_path: "/mnt/share1", client_share_path: "/mnt/client1", client_mount_options: "nosuid,rw,sync,hard,intr" }</code></li> <li>- <code>{ server_ip: xx.xx.xx.xx, server_share_path: "/mnt/share2", client_share_path: "/mnt/client2", client_mount_options: "nosuid,rw,sync,hard,intr" }</code></li> <li>- <code>{ server_ip: yy.yy.yy.yy, server_share_path: "/mnt/share3", client_share_path: "/mnt/client3", client_mount_options: "nosuid,rw,sync,hard,intr" }</code></li> </ul> </li> </ol> <p><b>Default value:</b> <code>{ server_ip: , server_share_path: , client_share_path: , client_mount_options: }</code></p> |
| <b>beegfs_support</b><br>boolean Optional      | <p>This variable is used to install beegfs-client on compute and manager nodes</p> <p>Choices:</p> <ul style="list-style-type: none"> <li>• false &lt;- Default</li> <li>• true</li> </ul>  |
| <b>beegfs_rdma_support</b><br>boolean Optional | <p>This variable is used if user has RDMA-capable network hardware (e.g., InfiniBand)</p> <p>Choices:</p> <ul style="list-style-type: none"> <li>• false &lt;- Default</li> <li>• true</li> </ul>   |

**telemetry\_config.yml**



Table 9: Parameters

| Parameter   | Details   |
|---|---|
| <b>idrac_telemetry_support</b><br>boolean <sup>Page 59, 1</sup><br>Required | <ul style="list-style-type: none"> <li>Enables iDRAC telemetry support and visualizations.</li> <li><b>Values:</b></li> </ul> <pre>* false &lt;- Default</pre> <pre>* true</pre> <hr/> <p><b>Note:</b> When <code>idrac_telemetry_support</code> is true, <code>mysqldb_user</code>, <code>mysqldb_password</code> and <code>mysqldb_root_password</code> become mandatory.</p> <hr/>   |
| <b>omnia_telemetry_support</b><br>boolean <sup>Page 59, 1</sup><br>Required | <ul style="list-style-type: none"> <li>Starts or stops Omnia telemetry</li> <li>If <code>omnia_telemetry_support</code> is true, then at least one of <code>collect_regular_metrics</code> or <code>collect_health_check_metrics</code> or <code>collect_gpu_metrics</code> should be true, to collect metrics.</li> <li>If <code>omnia_telemetry_support</code> is false, telemetry acquisition will be stopped.</li> <li><b>Values:</b></li> </ul> <pre>* false &lt;- Default</pre> <pre>* true</pre>   |
| <b>visualization_support</b><br>boolean <sup>Page 59, 1</sup><br>Required   | <ul style="list-style-type: none"> <li>Enables visualizations.</li> <li><b>Values:</b></li> </ul> <pre>* false &lt;- Default</pre> <pre>* true</pre> <hr/> <p><b>Note:</b> When <code>visualization_support</code> is true, <code>grafana_username</code> and <code>grafana_password</code> become mandatory.</p> <hr/>   |
| <b>appliance_k8s_pod_net_cidr</b><br>string<br>Required                     | <ul style="list-style-type: none"> <li>Kubernetes pod network CIDR for appliance k8s network.</li> <li>Make sure this value does not overlap with any of the host networks.</li> <li><b>Default value:</b> "192.168.0.0/16"</li> </ul>  |
| <b>pod_external_ip_start_range</b><br>string<br>Required                    | <ul style="list-style-type: none"> <li>The start of the range that will be used by Load-balancer for assigning IPs to K8s services in admin NIC subnet configured on the control plane.</li> <li>The first and second octets (x,y) are not used/validated by Omnia. These values are internally calculated based on the value of <code>admin_nic_subnet</code> in <code>input/provision_config.yml</code>.</li> <li>If <code>pod_external_ip_start_range:</code> "x.y.240.100" and <code>pod_external_ip_end_range:</code> "x.y.240.105" and</li> </ul> |
| 2.4. Configuring the cluster  | <ul style="list-style-type: none"> <li>If <code>admin_nic_subnet</code> provided in <code>provision_config.yml</code> is 10.5.0.0, <code>pod_external_ip_start_range</code> will be 10.5.240.100 and</li> </ul>   |

Click [here](#) for more information on [FreeIPA](#), [LDAP](#), [Telemetry](#), [BeeGFS](#) or, [NFS](#).

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

## 2.4.2 Before you build clusters

- [Ensure that all cluster nodes are up and running](#).
- Verify that all inventory files are updated.
- If the cluster requires more than 10 kubernetes nodes, use a docker enterprise account to avoid docker pull limits.
- Verify that all nodes are assigned a group. Use the [inventory](#) as a reference. The inventory file is case-sensitive. Follow the casing provided in the sample file link.
  - The manager group should have exactly 1 manager node.
  - The compute group should have at least 1 node.
  - The login group is optional. If present, it should have exactly 1 node.

---

### Note:

- The inventory file accepts both IPs and FQDNs as long as they can be resolved by DNS.
- In a multi-node setup, IP's cannot be repeated in the manager or compute groups. That is, don't include the manager node IP address in the compute group. In a single node setup, the compute node and the manager node must be the same.

- 
- Users should also ensure that all repos are available on the cluster nodes running RHEL.
  - Nodes provisioned using the Omnia provision tool do not require a RedHat subscription to run `omnia.yml` on RHEL cluster nodes.
  - For RHEL cluster nodes not provisioned by Omnia, ensure that RedHat subscription is enabled on all target nodes. Every target node will require a RedHat subscription.

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

## 2.4.3 Building clusters

1. In the `input/omnia_config.yml`, `input/security_config.yml`, `input/telemetry_config.yml` and [optional] `input/storage_config.yml` files, provide the [required details](#).

---

### Note:

- Use the parameter `scheduler_type` in `input/omnia_config.yml` to customize what schedulers are installed in the cluster.
- Without the login node, Slurm jobs can be scheduled only through the manager node.

- 
2. Create an inventory file in the `omnia` folder. Check out the [sample inventory for more information](#).

- `[Manager]` group should contain the manager node IP address.<sup>1</sup>
- `[compute]` group should contain all compute node IP addresses.

---

<sup>1</sup> In a multi-node setup, IP's cannot be repeated in the manager or compute groups. That is, don't include the manager node IP address in the compute group. In a single node setup, the compute node and the manager node must be the same.

- `[login]` group should contain the login node IP.

### Hostname requirements

- The Hostname should not contain the following characters: , (comma), . (period) or \_ (underscore). However, the **domain name** is allowed commas and periods.
- The Hostname cannot start or end with a hyphen (-).
- No upper case characters are allowed in the hostname.
- The hostname cannot start with a number.
- The hostname and the domain name (that is: `hostname00000x.domain.xxx`) cumulatively cannot exceed 64 characters. For example, if the `node_name` provided in `input/provision_config.yml` is 'node', and the `domain_name` provided is 'omnia.test', Omnia will set the hostname of a target cluster node to 'node00001.omnia.test'. Omnia appends 6 digits to the hostname to individually name each target node.

### Note:

- RedHat nodes that are not configured by Omnia need to have a valid subscription. To set up a subscription, [click here](#).
- Omnia creates a log file which is available at: `/var/log/omnia.log`.
- If only Slurm is being installed on the cluster, docker credentials are not required.

3. `omnia.yml` is a wrapper playbook comprising of:

- `security.yml`: This playbook sets up centralized authentication (LDAP/FreeIPA) on the cluster. For more information, [click here](#).
- `scheduler.yml`: This playbook sets up job schedulers (Slurm or Kubernetes) on the cluster.
- `storage.yml`: This playbook sets up storage tools like [BeeGFS](#) and [NFS](#).
- `telemetry.yml`: This playbook sets up [Omnia telemetry and/or iDRAC telemetry](#). It also installs [Grafana](#) and [Loki](#) as Kubernetes pods.


To run `omnia.yml`:

```
ansible-playbook omnia.yml -i inventory
```

### Note:

- To visualize the cluster (Slurm/Kubernetes) metrics on Grafana (On the control plane) during the run of `omnia.yml`, add the parameters `grafana_username` and `grafana_password` (That is `ansible-playbook omnia.yml -i inventory -e grafana_username="" -e grafana_password=""`).
- Having the same node in the manager and login groups in the inventory is not recommended by Omnia.
- If you want to view or edit the `omnia_config.yml` file, run the following command:
  - `ansible-vault view omnia_config.yml --vault-password-file .omnia_vault_key` – To view the file.
  - `ansible-vault edit omnia_config.yml --vault-password-file .omnia_vault_key` – To edit the file.
- Use the `ansible-vault view` or `edit` commands and not the `ansible-vault decrypt` or `encrypt` commands. If you have used the `ansible-vault decrypt` or `encrypt` commands, provide 644 permission to the parameter files.

## Setting up a shared home directory



Users wanting to set up a shared home directory for the cluster can do it in one of two ways:

- **Using the head node as an NFS host:** Set `enable_omnia_nfs` (`input/omnia_config.yml`) to `true` and provide a share path which will be configured on all nodes in `omnia_usrhome_share` (`input/omnia_config.yml`). During the execution of `omnia.yml`, the NFS share will be set up for access by all cluster nodes.
- **Using an external filesystem:** Configure the external file storage using `storage.yml`. Set `enable_omnia_nfs` (`input/omnia_config.yml`) to `false` and provide the external share path in `omnia_usrhome_share` (`input/omnia_config.yml`). Run `omnia.yml` to configure access to the external share for deployments.

### Slurm job based user access

To ensure security while running jobs on the cluster, users can be assigned permissions to access cluster nodes only while their jobs are running. To enable the feature:

```
cd scheduler
ansible-playbook job_based_user_access.yml -i inventory
```

#### Note:

- The inventory queried in the above command is to be created by the user prior to running `omnia.yml` as `scheduler.yml` is invoked by `omnia.yml`
- Only users added to the 'slurm' group can execute slurm jobs. To add users to the group, use the command: `usermod -a -G slurm <username>`.

### Running Slurm MPI jobs on clusters

To enhance the productivity of the cluster, Slurm allows users to run jobs in a parallel-computing architecture. This is used to efficiently utilize all available computing resources. [Click here for more information.](#)

#### Note:

- Omnia does not install MPI packages by default. Users hoping to leverage the Slurm-based MPI execution feature are required to install the relevant packages from a source of their choosing. For information on setting up Intel OneAPI on the cluster, [click here](#).
  - Ensure there is an NFS node on which to host slurm scripts to run.
  - Running jobs as individual users (and not as root) requires that passwordSSH be enabled between cluster nodes for the user.
- 

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

## 2.4.4 Centralized authentication systems

To enable centralized authentication in the cluster, Omnia installs either:

- FreeIPA
  - LDAP Client
- 

### Note:

- Nodes provisioned using the Omnia provision tool do not require a RedHat subscription to run `security.yml` on RHEL cluster nodes.
  - For RHEL cluster nodes not provisioned by Omnia, ensure that RedHat subscription is enabled on all cluster nodes. Every cluster node will require a RedHat subscription.
- 

### Using FreeIPA

Enter the following parameters in `input/security_config.yml`.

**Caution:** Do not remove or comment any lines in the `input/security_config.yml` file.

| Parameter  | Details  |
|--|--|
| <b>freeipa_required</b><br>boolean Required          | Boolean indicating whether FreeIPA is required or not.<br>Choices: <ul style="list-style-type: none"> <li>• true &lt;- Default</li> <li>• false</li> </ul>   |
| <b>realm_name</b><br>string Required                 | Sets the intended realm name.<br><b>Default value:</b> OMNIA.TEST  |
| <b>directory_manager_password</b><br>string Required | <ul style="list-style-type: none"> <li>• <b>Password authenticating admin level access to the Directory for system management tasks.</b> <ul style="list-style-type: none"> <li>– It will be added to the instance of directory server created for IPA.</li> <li>– Required Length: 8 characters.</li> <li>– The password must not contain -, ‘,’</li> </ul> </li> </ul> |
| <b>kerberos_admin_password</b><br>string Required    | “admin” user password for the IPA server on RockyOS.   |
| <b>domain_name</b><br>string Required                | Sets the intended domain name<br><b>Default value:</b> omnia.test  |

**Note:** The input/security\_config.yml file is encrypted on the first run of security.yml and omnia.yml:

To view the encrypted parameters:

```
ansible-vault view security_config.yml --vault-password-file .security_vault.  
↪key
```

To edit the encrypted parameters:

```
ansible-vault edit security_config.yml --vault-password-file .security_vault.  
↪key
```

If a subsequent run of security.yml or omnia.yml, all configuration files that have been encrypted by the playbook will be unencrypted.

Omnia installs a FreeIPA server on the manager node and FreeIPA clients on the cluster and login node using one of the below commands:

```
ansible-playbook security.yml -i inventory  
  
ansible-playbook omnia.yml -i inventory
```

Where inventory follows the format defined under inventory file in the provided [Sample Files](#). The inventory file is case-sensitive. Follow the casing provided in the sample file link.

The omnia.yml playbook installs Slurm, BeeFGS Client, NFS Client in addition to freeIPA.

**Note:**

- Omnia does not create any accounts (HPC users) on FreeIPA. To create a user, check out *FreeIPA documentation*.
- Alternatively, use the below commands with admin credentials on the login/head node:

```
kinit admin (When prompted, provide kerberos_admin_password as entered in security_
↪config.yml)
ipa user-add FirstName_LastName --first=FirstName --last=LastName --password --
↪homedir=/home/omnia-share/FirstName_LastName --shell /bin/bash
```

For example: `ipa user-add FirstName_LastName --first=FirstName --last=LastName --password --homedir=/home/omnia-share/FirstName_LastName --shell /bin/bash`

After the new user account logs in for the first time, you will be prompted to change the password to the account.

**Setting up Passwordless SSH for FreeIPA**

Once user accounts are created, admins can enable passwordless SSH for users to run HPC jobs on the cluster nodes.

**Note:** Once user accounts are created on FreeIPA, use the accounts to login to the cluster nodes to reset the password and create a corresponding home directory.

To customize your setup of passwordless ssh, input parameters in `input/passwordless_ssh_config.yml`.

| Parameter                                       | Details   |
|---|---|
| <b>user_name</b><br>string Required             | <b>The list of users that requires passwordless SSH. Separate the list of users using a comma.</b><br>Eg: user1,user2,user3   |
| <b>authentication_type</b><br>string Required   | Indicates whether LDAP or FreeIPA is in use on the cluster.<br>Choices: <ul style="list-style-type: none"> <li>• freeipa &lt;- Default</li> <li>• ldap</li> </ul>   |
| <b>freeipa_user_home_dir</b><br>string Required | <ul style="list-style-type: none"> <li>• <b>This variable accepts the user home directory path for freeipa configuration.</b> <ul style="list-style-type: none"> <li>– If nfs mount is created for user home, make sure you provide the freeipa users mount home directory path.</li> </ul> </li> </ul> <b>Default value:</b> "/home/omnia-share" |

Use the below command to enable passwordless SSH:

```
ansible-playbook user_passwordless_ssh.yml -i inventory
```

Where inventory follows the format defined under inventory file in the provided [Sample Files](#). The inventory file is case-sensitive. Follow the casing provided in the sample file link.

**Caution:** Do not run ssh-keygen commands after passwordless SSH is set up on the nodes.

### **Using LDAP client**

To add the cluster to an external LDAP server, Omnia enables the installation of LDAP client on the manager, compute and login nodes.

To customize your LDAP client installation, input parameters in `input/security_config.yml`



| Parameter                                      | Details  |
|--|--|
| <b>ldap_required</b><br>boolean Required       | Boolean indicating whether LDAP is required or not.<br>Choices: <ul style="list-style-type: none"> <li>• true &lt;- Default</li> <li>• false</li> </ul>  |
| <b>domain_name</b><br>string Required          | Sets the intended domain name<br><b>Default value:</b> omnia.test  |
| <b>ldap_server_ip</b><br>string Optional       | LDAP server IP. Required if ldap_required is true. There should be an explicit LDAP server running on this IP.   |
| <b>ldap_connection_type</b><br>string Required | <ul style="list-style-type: none"> <li>• For a TLS connection, provide a valid certification path.</li> <li>• For an SSL connection, ensure port 636 is open.</li> </ul> Choices: <ul style="list-style-type: none"> <li>– TLS &lt;- Default</li> <li>– SSL</li> </ul>   |
| <b>ldap_ca_cert_path</b><br>string Required    | <ul style="list-style-type: none"> <li>• This variable accepts Server Certificate Path.</li> <li>• Make sure certificate is present in the path provided.</li> <li>• The certificate should have .pem or .crt extension.</li> <li>• This variable is mandatory if connection type is TLS.</li> </ul> <b>Default value:</b> /etc/openldap/certs/omnialdap.pem               |
| <b>user_home_dir</b><br>string Required        | <ul style="list-style-type: none"> <li>• This variable accepts the user home directory path for LDAP configuration.</li> <li>• If nfs mount is created for user home, make sure you provide the freeipa users mount home directory path.</li> </ul> <b>Default value:</b> "/home/omnia-share"  |
| <b>ldap_bind_username</b><br>string Required   | <ul style="list-style-type: none"> <li>• If LDAP server is configured with bind dn then bind dn user to be provided.</li> <li>• If this value is not provided (when bind is configured in server) then ldap authentication fails.</li> <li>• Omnia does not validate this input.</li> <li>• Ensure that it is valid and proper.</li> </ul> <b>Default value:</b> admin     |
| <b>ldap_bind_password</b><br>string Required   | <ul style="list-style-type: none"> <li>• If LDAP server is configured with bind dn then bind dn password to be provided.</li> <li>• If this value is not provided (when bind is configured in server) then ldap authentication fails.</li> <li>• Omnia does not validate this input.</li> <li>• Ensure that it is valid and proper.</li> </ul> <b>Default value:</b> admin |
| <b>2.4. Configuring the cluster</b>            | <b>Default value:</b> admin <b>73</b>  |

---

**Note:** Omnia does not create any accounts (HPC users) on LDAP. To create a user, check out [LDAP documentation](#).

---

### Setting up Passwordless SSH for LDAP

Once user accounts are created, admins can enable passwordless SSH for users to run HPC jobs on the cluster nodes.

---

**Note:**

- Ensure that the control plane can reach the designated LDAP server.
  - If `enable_omnia_nfs` is true in `input/omnia_config.yml`, follow the below steps to configure an NFS share on your LDAP server:
    - **From the manager node:**
      1. Add the LDAP server IP address to `/etc/exports`.
      2. Run `exportfs -ra` to enable the NFS configuration.
    - **From the LDAP server:**
      1. Add the required `fstab` entries in `/etc/fstab`. (The corresponding entry will be available on the cluster nodes in `/etc/fstab`)
      2. Mount the NFS share using `mount manager_ip: /home/omnia-share /home/omnia-share`.
  - If `enable_omnia_nfs` is false in `input/omnia_config.yml`, ensure the user-configured NFS share is mounted on the LDAP server.
- 

To customize your setup of passwordless ssh, input parameters in `input/passwordless_ssh_config.yml`

| Parameter  | Details   |
|--|---|
| <b>user_name</b><br>string Required                | <b>The list of users that requires passwordless SSH. Separate the list of users using a comma.</b><br>Eg: user1,user2,user3   |
| <b>authentication_type</b><br>string Required      | Indicates whether LDAP or FreeIPA is in use on the cluster.<br><br>Choices: <ul style="list-style-type: none"><li>• <code>freeipa</code> &lt;- Default</li><li>• <code>ldap</code></li></ul>  |
| <b>ldap_organizational_unit</b><br>string Required | <ul style="list-style-type: none"><li>• Distinguished name i.e dn in ldap is used to identify an entity in a LDAP.</li><li>• This variable includes the organizational unit (ou) which is used to identifies user in the LDAP.</li><li>• Only provide ou details i.e ou=people, as domain name and userid is accepted already.</li></ul> <b>Default value:</b> people |

Use the below command to enable passwordless SSH:

```
ansible-playbook user_passwordless_ssh.yml -i inventory
```

Where inventory follows the format defined under inventory file.

```
[manager]
10.5.0.101

[compute]
10.5.0.102
10.5.0.103


[ldap_server]
10.5.0.105
```

**Caution:** Do not run ssh-keygen commands after passwordless SSH is set up on the nodes.

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

## 2.4.5 BeeGFS bolt on

BeeGFS is a hardware-independent POSIX parallel file system (a.k.a. Software-defined Parallel Storage) developed with a strong focus on performance and designed for ease of use, simple installation, and management.



### Pre Requisites before installing BeeGFS client

- Ensure that the BeeGFS server is set up using the [linked steps](#).
- Ensure that a `connAuthFile` is configured on the server as explained [here](#)

**Caution:** Configuring a `connAuthFile` is now mandatory. Services will no longer start if a `connAuthFile` is not configured

- Ensure that the following ports are open for TCP and UDP connectivity:

| Port | Service                          |
|------|----------------------------------|
| 8008 | Management service (beegfs-mgmt) |
| 8003 | Storage service (beegfs-storage) |
| 8004 | Client service (beegfs-client)   |
| 8005 | Metadata service (beegfs-meta)   |
| 8006 | Helper service (beegfs-helper)   |

To open the ports required, use the following steps:

1. `firewall-cmd --permanent --zone=public --add-port=<port number>/tcp`
2. `firewall-cmd --permanent --zone=public --add-port=<port number>/udp`
3. `firewall-cmd --reload`
4. `systemctl status firewalld`

- Ensure that the nodes in the inventory have been assigned **only** these roles: manager and compute.
- Nodes provisioned using the Omnia provision tool do not require a RedHat subscription to set up BeeGFS on RHEL cluster nodes.
- For RHEL cluster nodes not provisioned by Omnia, ensure that RedHat subscription is enabled on all cluster nodes. Every cluster node will require a RedHat subscription.

---

**Note:** BeeGFS services over RDMA is only supported on RHEL 8.3 and above due to limitations on BeeGFS. When setting up your cluster with RDMA support, check the BeeGFS documentation to provide appropriate values in `input/storage_config.yml`.

---

- If the cluster runs Rocky, ensure that versions running are compatible by checking our [support matrix](#).

### Installing the BeeGFS client via Omnia

After the required parameters are filled in `input/storage_config.yml`, Omnia installs BeeGFS on manager and compute nodes while executing the `omnia.yml` playbook.

**Caution:** Do not remove or comment any lines in the `input/storage_config.yml` file.

| Name  | Details   |
|---|---|
| <b>beegfs_support</b><br>boolean Optional                 | This variable is used to install beegfs-client on compute and manager nodes<br>Choices: <ul style="list-style-type: none"> <li>• false &lt;- Default</li> <li>• true</li> </ul>   |
| <b>beegfs_rdma_support</b><br>boolean Optional            | This variable is used if user has RDMA-capable network hardware (e.g., InfiniBand)<br>Choices: <ul style="list-style-type: none"> <li>• false &lt;- Default</li> <li>• true</li> </ul>  |
| <b>beegfs_ofed_kernel_modules_path</b><br>string Optional | The path where separate OFED kernel modules are installed.<br><b>Default value:</b> "/usr/src/ofa_kernel/default/include"   |
| <b>beegfs_mgmt_server</b><br>string Required              | BeeGFS management server IP. Note: The provided IP should have an explicit BeeGFS management server running .   |
| <b>beegfs_mounts</b><br>string Optional                   | <b>Beegfs-client file system mount location. If storage.yml is being used to change the BeeGFS mounts location, set beegfs_unmount_client to true</b><br><b>Default value:</b> "/mnt/beegfs"  |
| <b>beegfs_unmount_client</b><br>boolean Optional          | Changing this value to true will unmount running instance of BeeGFS client and should only be used when decommissioning BeeGFS, changing the mount location or changing the BeeGFS version.<br>Choices: <ul style="list-style-type: none"> <li>• false &lt;- Default</li> <li>• true</li> </ul> |
| <b>beegfs_client_version</b><br>string Optional           | Beegfs client version needed on compute and manager nodes.<br><b>Default value:</b> 7.2.6   |
| <b>beegfs_version_change</b><br>boolean Optional          | Use this variable to change the BeeGFS version on the target nodes.<br>Choices: <ul style="list-style-type: none"> <li>• false &lt;- Default</li> <li>• true</li> </ul>   |
| <b>beegfs_secret_storage_filepath</b><br>string Optional  | <ul style="list-style-type: none"> <li>• The filepath (including the filename) where the connauthfile is placed.</li> <li>• Recommended for Beegfs version &gt;= 7.2.7.</li> </ul> <b>Default value:</b> /home/connauthfile   |

**Note:**

- BeeGFS client-server communication can take place over TCP or RDMA. If RDMA support is required, set `beegfs_rdma_support` should be set to true. Also, OFED should be installed on all cluster nodes.
- For BeeGFS communication happening over RDMA, the `beegfs_mgmt_server` should be provided with the Infiniband IP of the management server.
- The parameter `inventory` refers to the [inventory file](#) listing manager, login and compute nodes.)

If `input/storage_config.yml` is populated before running `omnia.yml`, BeeGFS client will be set up during the run of `omnia.yml`.

If `omnia.yml` is not leveraged to set up BeeGFS, run the `storage.yml` playbook :

```
cd storage
ansible-playbook storage.yml -i inventory
```

**Note:** Once BeeGFS is successfully set up, set `enable_omnia_nfs` (`input/omnia_config.yml`) to false and `omnia_usrhome_share` (`input/omnia_config.yml`) to an accessible share path in BeeGFS to use the path across the cluster for deployments.

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

## 2.4.6 NFS bolt on

- Ensure that an external NFS server is set up using the [linked steps](#) alternatively, `nfs_sas.yml` can be leveraged. NFS clients are mounted using the external NFS server's IP.
- Fill out the `nfs_client_params` variable in the `input/storage_config.yml` file in JSON format using the samples provided below.
- This role runs on manager, compute and login nodes.
- Make sure that `/etc/exports` on the NFS server is populated with the same paths listed as `server_share_path` in the `nfs_client_params` in `input/storage_config.yml`.
- Post configuration, enable the following services (using this command: `firewall-cmd --permanent --add-service=<service name>`) and then reload the firewall (using this command: `firewall-cmd --reload`).
  - `nfs`
  - `rpc-bind`
  - `mountd`
- Omnia supports all NFS mount options. Without user input, the default mount options are `no-suid,rw,sync,hard,intr`.
- The fields listed in `nfs_client_params` are:
  - `server_ip`: IP of NFS server
  - `server_share_path`: Folder on which NFS server mounted
  - `client_share_path`: Target directory for the NFS mount on the client. If left empty, respective `server_share_path` value will be taken for `client_share_path`.

- `client_mount_options`: The mount options when mounting the NFS export on the client. Default value: `nosuid,rw,sync,hard,intr`. For a list of mount options, [click here](#).
- Nodes provisioned using the Omnia provision tool do not require a RedHat subscription to set up NFS on RHEL target nodes.
- For RHEL target nodes not provisioned by Omnia, ensure that RedHat subscription is enabled on all target nodes. Every target node will require a RedHat subscription.
- There are 3 ways to configure the feature:
  1. **Single NFS node** : A single NFS filesystem is mounted from a single NFS server. The value of `nfs_client_params` would be:

```
- { server_ip: 10.5.0.101, server_share_path: "/mnt/share", client_share_path:
  ↪ "/mnt/client", client_mount_options: "nosuid,rw,sync,hard,intr" }
```

2. **Multiple Mount NFS Filesystem**: Multiple filesystems are mounted from a single NFS server. The value of `nfs_client_params` would be:

```
- { server_ip: 10.5.0.101, server_share_path: "/mnt/share1", client_share_path:
  ↪ "/mnt/client1", client_mount_options: "nosuid,rw,sync,hard,intr" }
- { server_ip: 10.5.0.101, server_share_path: "/mnt/share2", client_share_path:
  ↪ "/mnt/client2", client_mount_options: "nosuid,rw,sync,hard,intr" }
```

3. **Multiple NFS Filesystems**: Multiple filesystems are mounted from multiple NFS servers. The value of `nfs_client_params` would be:

```
- { server_ip: 10.5.0.101, server_share_path: "/mnt/server1", client_share_
  ↪ path: "/mnt/client1", client_mount_options: "nosuid,rw,sync,hard,intr" }
- { server_ip: 10.5.0.102, server_share_path: "/mnt/server2", client_share_
  ↪ path: "/mnt/client2", client_mount_options: "nosuid,rw,sync,hard,intr" }
- { server_ip: 10.5.0.103, server_share_path: "/mnt/server3", client_share_
  ↪ path: "/mnt/client3", client_mount_options: "nosuid,rw,sync,hard,intr" }
```

**Caution:** After an NFS client is configured, if the NFS server is rebooted, the client may not be able to reach the server. In those cases, restart the NFS services on the server using the below command:

```
systemctl disable nfs-server
systemctl enable nfs-server
systemctl restart nfs-server
```

If `omnia.yml` is not leveraged to set up NFS, run the `storage.yml` playbook :

```
cd storage
ansible-playbook storage.yml -i inventory
```

**Note:** Once NFS is successfully set up, set `enable_omnia_nfs` (`input/omnia_config.yml`) to false and `omnia_usrhome_share` (`input/omnia_config.yml`) to an accessible share path in the NFS share to use the path across the cluster for deployments.

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

## 2.4.7 Updating kernels

### Pre requisites:

- Verify that the BaseOS repo version on the control plane matches the cluster kernel version for the cluster nodes.

```
dnf repolist
```

- Verify that the cluster nodes are in the booted state. For more information, [click here](#).

### Download and install kernel updates to cluster nodes

1. Download all required kernel RPM update files to the `/install/post/otherpkgs/<Provision OS.Version>/x86_64/custom_software/Packages/` folder.

For example:

```
cd /install/post/otherpkgs/rhels8.6.0/x86_64/custom_software/Packages/  
dnf download kernel* --resolve -y
```

2. Once downloaded, make sure there are `>=20~` rpm packages in the `/install/post/otherpkgs/rhels8.6.0/x86_64/custom_software/Packages/` directory.
3. Inside the `/install/post/otherpkgs/<Provision OS.Version>/x86_64/custom_software` folder, create a file named `update.otherpkgs.pkglist` with the contents:

```
custom_software/kernel*
```

4. Go to `utils/os_package_update` and edit `package_update_config.yml`. For more information on the input parameters, [click here](#).
5. Run `package_update.yml` using: `ansible-playbook package_update.yml`
6. After execution is completed, verify that kernel packages are on the nodes using: `rpm -qa | grep kernel`

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

## 2.5 Adding new nodes

### Provisioning the new node

While adding a new node to the cluster, users can modify the following:

- The operating system
- CUDA
- OFED

A new node can be added using the following ways:

- When the discovery mechanism is mapping:
  - Update the existing mapping file by appending the new entry (without the disrupting the older entries) or provide a new mapping file by pointing `pxe_mapping_file_path` in `provision_config.yml` to the new location.



---

**Note:** When re-running `provision.yml` with a new mapping file, ensure that existing IPs from the current mapping file are not provided in the new mapping file. Any IP overlap between mapping files will result in PXE failure. This can only be resolved by running the [Clean Up script](#) followed by `provision.yml`.

---

- Run `provision.yml`:

```
cd provision
ansible-playbook provision.yml
```

- Manually PXE boot the target servers after the `provision.yml` playbook is executed and the target node lists as **booted** in the [nodeinfo table](#)

- When the discovery mechanism is `bmc`:

- Run `provision.yml` once the node has joined the cluster using an IP that exists within the provided range.

```
cd provision
ansible-playbook provision.yml
```

- When the discovery mechanism is `switch-based`:

- Edit or append JSON list stored in `switch-based-details` in `input/provision_config.yml`.

---

**Note:**

- All ports residing on the same switch should be listed in the same JSON list element.
  - Ports configured via Omnia should not be removed from `switch-based-details` in `input/provision_config.yml`.
- 

- Run `provision.yml`.

```
cd provision
ansible-playbook provision.yml
```

- Manually PXE boot the target servers after the `provision.yml` playbook is executed and the target node lists as **booted** in the [nodeinfo table](#)

- When the discovery mechanism is `snmpwalk`:

- Run `provision.yml` after the switch has discovered the new node.

```
cd provision
ansible-playbook provision.yml
```

- Manually PXE boot the target servers after the `provision.yml` playbook is executed and the target node lists as **booted** in the [Omnia nodeinfo table](#).

Alternatively, if a new node is to be added with no change in configuration, run the following commands:

```
cd provision
ansible-playbook discovery_provision.yml
```

Verify that the node has been provisioned successfully by [checking the Omnia nodeinfo table](#).

### Adding the new node to the cluster

1. Insert the new IPs in the existing inventory file following the below example.

*Existing inventory*

```
[manager]
10.5.0.101

[compute]
10.5.0.102
10.5.0.103

[login]
10.5.0.104
```

*Updated inventory with the new node information*

```
[manager]
10.5.0.101

[compute]
10.5.0.102
10.5.0.103
10.5.0.105
10.5.0.106

[login]
10.5.0.104
```

In the above example, nodes 10.5.0.105 and 10.5.0.106 have been added to the cluster as a compute nodes.

---

#### Note:

- Do not change the manager node in the existing inventory. Simply add the new node information in the compute group.
- Only the `scheduler_type` in `input/omnia_config.yml` and the variables in `input/storage_config.yml` can be updated while re-running `omnia.yml` to add the new node. All other variables in the files `input/omnia_config.yml` and `input/security_config.yml` must be unedited.

- 
3. To install [security](#), [job scheduler](#) and storage tools ([NFS](#), [BeeGFS](#)) on the node, run `omnia.yml`:

```
ansible-playbook omnia.yml -i inventory
```

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

## 2.6 Re-provisioning the cluster

In the event that an existing Omnia cluster needs a different OS version or a fresh installation, the cluster can be re-provisioned. While re-provisioning the cluster, users can modify the following:

- The operating system
- CUDA
- OFED

Omnia can re-provision the cluster by running the following command:

```
cd provision
ansible-playbook provision.yml -i inventory
```

Alternatively, if a re-deployment with no modifications are required

```
cd provision
ansible-playbook discovery_provision.yml -i inventory
```

Where the inventory contains a list of host IPs (Sourced from the [nodeinfo table](#)) as shown below:

```
10.5.0.101
10.5.0.102
```

---

### Note:

- The host IPs passed in the inventory should be assigned by Omnia. They will not be changed during the re-provisioning.
  - If the nodes were discovered via snmpwalk or mapping, users will be required to manually reboot target nodes.
  - Do not include groups like *manager*, *compute* and *login* in the passed inventory.
- 

### Setting up the cluster

1. Insert the new IPs (only if a new node is to be added) and/or move nodes between groups in the existing inventory file following the below example.

*Existing inventory*

```
[manager]
10.5.0.101

[compute]
10.5.0.102
10.5.0.103

[login]
10.5.0.104
```

*Updated inventory with the new node information*

```
[manager]
10.5.0.102
```

(continues on next page)

(continued from previous page)

```
[compute]
10.5.0.101
10.5.0.103
10.5.0.104
10.5.0.106

[login]
10.5.0.105
```

In the above example, the compute node: 10.5.0.102 has been moved to manager, a new login node: 10.5.0.105 has been set up, and the node: 10.5.0.106 has been added to the cluster as a compute node.

---

**Note:** To change the configuration of the cluster, the files `input/omnia_config.yml`, `input/security_config.yml` and `input/storage_config.yml` can be updated before running `omnia.yml`.

---

3. To install `security`, `job scheduler` and storage tools (NFS, BeeGFS) on the node, run `omnia.yml`:

```
ansible-playbook omnia.yml -i inventory
```

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

## 2.7 Configuring switches

### 2.7.1 Configuring infiniband switches

Depending on the number of ports available on your Infiniband switch, they can be classified into:

- EDR Switches (36 ports)
- HDR Switches (40 ports)
- NDR Switches (32 ports)

Input the configuration variables into the `network/infiniband_edr_input.yml`, `network/infiniband_hdr_input.yml` or `network/infiniband_ndr_input.yml` as appropriate:

**Caution:** Do not remove or comment any lines in the `network/infiniband_edr_input.yml`, `network/infiniband_hdr_input.yml` or `network/infiniband_ndr_input.yml` file.

| Parameters   | Details  |
|--|--|
| <b>enable_split_port</b><br>boolean <sup>1</sup> Required                | Indicates whether ports are to be split.<br>Choices: <ul style="list-style-type: none"> <li>• false &lt;- default</li> <li>• true</li> </ul>   |
| <b>ib_split_ports</b><br>string Optional                                 | <ul style="list-style-type: none"> <li>• Stores the split configuration of the ports.</li> <li>• For EDR and HDR switches, the accepted formats are : comma-separated (EX: “1,2”), ranges (EX: “1-10”), comma-separated ranges (EX: “1,2,3-8,9,10-12”)</li> <li>• For NDR switches, the accepted format is: 2/1, 2/2, 3/1</li> </ul> <hr/> <p><b>Note:</b> The port prefix IB1 can be ignored when setting this value.</p> <hr/> |
| <b>snmp_community_name</b><br>string Optional                            | <p>The “SNMP community string” is like a user ID or password that allows access to a router’s or other device’s statistics.</p> <p><b>Default values:</b> public</p>   |
| <b>cache_directory</b><br>string Optional                                | Cache location used by OpenSM  |
| <b>log_directory</b><br>string Optional                                  | The directory where temporary files of opensm are stored. Can be set to the default directory or enter a directory path to store temporary files.  |
| <b>mellanox_switch_config</b><br>string Optional                         | <p>List of configuration lines to apply to the switch. # Example:</p> <p><b># mellanox_switch_config:</b><br/># - Command 1 # - Command 2</p> <p>By default, the list is empty.</p>  |
| <b>ib 1/(1-xx) config</b><br>string Optional                             | <p>Indicates the required state of ports 1-xx (depending on the value of 1/x).</p> <p><b>Default values:</b> "no shutdown"</p>   |
| <b>save_changes_to_startup</b><br>boolean <sup>Page 85, 1</sup> Optional | <p><b>Indicates whether the switch configuration is to persist across reboots.</b></p> <p>Choices:</p> <ul style="list-style-type: none"> <li>• false &lt;- default</li> <li>• true</li> </ul>   |

### Before you run the playbook

Before running `network/infiniband_switch_config.yml`, ensure that SSL Secure Cookies are disabled. Also, HTTP and JSON Gateway need to be enabled on your switch. This can be verified by running:

<sup>1</sup> Boolean parameters do not need to be passed with double or single quotes.

```
show web (To check if SSL Secure Cookies is disabled and HTTP is enabled)
show json-gw (To check if JSON Gateway is enabled)
```

In case any of these services are not in the state required, run:

```
no web https ssl secure-cookie enable (To disable SSL Secure Cookies)
web http enable (To enable the HTTP gateway)
json-gw enable (To enable the JSON gateway)
```

When connecting to a new or factory reset switch, the configuration wizard requests to execute an initial configuration:  
(Recommended) If the user enters ‘no’, they still have to provide the admin and monitor passwords.

If the user enters ‘yes’, they will also be prompted to enter the hostname for the switch, DHCP details, IPv6 details, etc.

---

**Note:**

- Currently, Omnia only supports the splitting of switch ports. Switch ports cannot be un-split using this script. For information on manually un-splitting ports, [click here](#).
  - When initializing a factory reset switch, the user needs to ensure DHCP is enabled and an IPv6 address is not assigned.
  - All ports intended for splitting need to be connected to the network before running the playbook.
  - The `ib_password` remains unchanged on switches that are in split-ready mode.
- 

**Running the playbook**

If `enable_split_port` is **true**, run:

```
cd network
ansible-playbook infiniband_switch_config.yml -i inventory -e ib_username="" -e ib_
↪password="" -e ib_admin_password="" -e ib_monitor_password="" -e ib_default_password=""
↪ -e ib_switch_type=""
```

If `enable_split_port` is **false**, run:

```
cd network
ansible-playbook infiniband_switch_config.yml -i inventory -e ib_username="" -e ib_
↪password="" -e ib_switch_type=""
```

- Where `ib_username` is the username used to authenticate into the switch.
  - Where `ib_password` is the password used to authenticate into the switch.
  - Where `ib_admin_password` is the intended password to authenticate into the switch after `infiniband_switch_config.yml` has run.
  - Where `ib_monitor_password` is the mandatory password required while running the initial configuration wizard on the Infiniband switch.
- 

**Note:**

- `ib_admin_password` and `ib_monitor_password` have the following constraints:
    - Passwords should contain 8-64 characters.
    - Passwords should be different from username.
-

- Passwords should be different from 5 previous passwords.
    - Passwords should contain at least one of each: Lowercase, uppercase and digits.
  - The inventory file should be a list of IPs separated by newlines. Check out the `switch_inventory` section in [Sample Files](#)
- 

- Where `ib_default_password` is the password used to authenticate into factory reset/fresh-install switches.
- Where `ib_switch_type` refers to the model of the switch: HDR/EDR/NDR

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

### 2.7.2 Configuring ethernet switches (S3 and S4 series)

- Edit the `network/ethernet_tor_input.yml` file for all S3\* and S4\* PowerSwitches such as S3048-ON, S4048T-ON, S4112F-ON, S4048-ON, S4048T-ON, S4112F-ON, S4112T-ON, and S4128F-ON.

|   |
|---|
| <b>Caution:</b> Do not remove or comment any lines in the <code>network/ethernet_tor_input.yml</code> file. |
|---|

| Name  | Details   |
|---|---|
| <b>os10_config</b><br>string Required                           | Global configurations for the switch.<br>Choices: <ul style="list-style-type: none"> <li>• <code>interface vlan1 &lt;- Default</code></li> <li>• <code>exit</code></li> </ul>   |
| <b>breakout_value</b><br>string Required                        | By default, all ports are configured in the 10g-4x breakout mode in which a QSFP28 or QSFP+ port is split into four 10G interfaces.<br>Choices: <ul style="list-style-type: none"> <li>• <code>10g-4x &lt;- Default</code></li> <li>• <code>25g-4x</code></li> <li>• <code>40g-1x</code></li> <li>• <code>50g-2x</code></li> <li>• <code>100g-1x</code></li> </ul>  |
| <b>snmp_trap_destination</b><br>string Optional                 | The trap destination IP address is the IP address of the SNMP Server where the trap will be sent. Ensure that the SNMP IP is valid.   |
| <b>snmp_community_string</b><br>string Optional                 | An SNMP community string is a means of accessing statistics stored within a router or other device.<br><b>Default values:</b> <code>public</code>   |
| <b>ethernet 1/1/(1-52) config</b><br>string Required            | By default: <ul style="list-style-type: none"> <li>• Port description is provided.</li> <li>• Each interface is set to “up” state.</li> <li>• The fanout/breakout mode for 1/1/1 to 1/1/52 is as per the value set in the <code>breakout_value</code> variable.</li> <li>• Update the individual interfaces of the Dell PowerSwitch S5232F-ON.</li> <li>• The interfaces are from ethernet 1/1/1 to ethernet 1/1/52. By default, the breakout mode is set for 1/1/1 to 1/1/52.</li> <li>• Note: The playbooks will fail if any invalid configurations are entered.</li> </ul> |
| <b>save_changes_to_startup</b><br>boolean <sup>1</sup> Required | Change it to “true” only when you are certain that the updated configurations and commands are valid.<br><b>WARNING:</b> When set to “true”, the startup configuration file is updated. If incorrect configurations or commands are entered, the Ethernet switches may not operate as expected.<br>Choices: <ul style="list-style-type: none"> <li>• <code>false &lt;- Default</code></li> <li>• <code>true</code></li> </ul>   |

- When initializing a factory reset switch, the user needs to ensure DHCP is enabled and an IPv6 address is not assigned.

### Running the playbook:

<sup>1</sup> Boolean parameters do not need to be passed with double or single quotes.



```
cd network
```

```
ansible-playbook ethernet_switch_config.yml -i inventory -e ethernet_switch_username="" -  
↪e ethernet_switch_password=""
```

- Where `ethernet_switch_username` is the username used to authenticate into the switch.
- The inventory file should be a list of IPs separated by newlines. Check out the `switch_inventory` section in [Sample Files](#)
- Where `ethernet_switch_password` is the password used to authenticate into the switch.

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

### 2.7.3 Configuring ethernet switches (S5 series)

- Edit the `network/ethernet_sseries_input.yml` file for all S5\* PowerSwitches such as S5232F-ON.

**Caution:** Do not remove or comment any lines in the `network/ethernet_sseries_input.yml` file.

| Name  | Details   |
|---|---|
| <b>os10_config</b><br>string Required                           | Global configurations for the switch.<br>Choices: <ul style="list-style-type: none"> <li>• <code>interface vlan1 &lt;- Default</code></li> <li>• <code>exit</code></li> </ul>   |
| <b>breakout_value</b><br>string Required                        | By default, all ports are configured in the 10g-4x breakout mode in which a QSFP28 or QSFP+ port is split into four 10G interfaces.<br>Choices: <ul style="list-style-type: none"> <li>• <code>10g-4x &lt;- Default</code></li> <li>• <code>25g-4x</code></li> <li>• <code>40g-1x</code></li> <li>• <code>50g-2x</code></li> <li>• <code>100g-1x</code></li> </ul>  |
| <b>snmp_trap_destination</b><br>string Optional                 | The trap destination IP address is the IP address of the SNMP Server where the trap will be sent. Ensure that the SNMP IP is valid.   |
| <b>snmp_community_string</b><br>string Optional                 | An SNMP community string is a means of accessing statistics stored within a router or other device.<br><b>Default values:</b> <code>public</code>   |
| <b>ethernet 1/1/(1-34) config</b><br>string Required            | By default: <ul style="list-style-type: none"> <li>• Port description is provided.</li> <li>• Each interface is set to “up” state.</li> <li>• The fanout/breakout mode for 1/1/1 to 1/1/31 is as per the value set in the <code>breakout_value</code> variable.</li> <li>• Update the individual interfaces of the Dell PowerSwitch S5232F-ON.</li> <li>• The interfaces are from ethernet 1/1/1 to ethernet 1/1/34. By default, the breakout mode is set for 1/1/1 to 1/1/34.</li> <li>• Note: The playbooks will fail if any invalid configurations are entered.</li> </ul> |
| <b>save_changes_to_startup</b><br>boolean <sup>1</sup> Required | Change it to “true” only when you are certain that the updated configurations and commands are valid.<br><b>WARNING:</b> When set to “true”, the startup configuration file is updated. If incorrect configurations or commands are entered, the Ethernet switches may not operate as expected.<br>Choices: <ul style="list-style-type: none"> <li>• <code>false &lt;- Default</code></li> <li>• <code>true</code></li> </ul>   |

- When initializing a factory reset switch, the user needs to ensure DHCP is enabled and an IPv6 address is not assigned.

<sup>1</sup> Boolean parameters do not need to be passed with double or single quotes.

---

**Note:** The `breakout_value` of a port can only be changed after un-splitting the port.

---

### Running the playbook:

```
cd network

ansible-playbook ethernet_switch_config.yml -i inventory -e ethernet_switch_username="" -
↪e ethernet_switch_password=""
```

- Where `ethernet_switch_username` is the username used to authenticate into the switch.
- The inventory file should be a list of IPs separated by newlines. Check out the `switch_inventory` section in [Sample Files](#)
- Where `ethernet_switch_password` is the password used to authenticate into the switch.

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

## 2.7.4 Configuring ethernet switches (Z series)

- Edit the `network/ethernet_zseries_input.yml` file for all Z series PowerSwitches such as Z9332F-ON, Z9262-ON and Z9264F-ON. The default configuration is written for Z9264F-ON.

**Caution:** Do not remove or comment any lines in the `network/ethernet_zseries_input.yml` file.

| Name  | Details   |
|---|---|
| <b>os10_config</b><br>string Required                           | Global configurations for the switch.<br>Choices: <ul style="list-style-type: none"> <li>• <code>interface vlan1 &lt;- Default</code></li> <li>• <code>exit</code></li> </ul>   |
| <b>breakout_value</b><br>string Required                        | By default, all ports are configured in the 10g-4x breakout mode in which a QSFP28 or QSFP+ port is split into four 10G interfaces.<br>Choices: <ul style="list-style-type: none"> <li>• <code>10g-4x &lt;- Default</code></li> <li>• <code>25g-4x</code></li> <li>• <code>40g-1x</code></li> <li>• <code>50g-2x</code></li> <li>• <code>100g-1x</code></li> </ul>  |
| <b>snmp_trap_destination</b><br>string Optional                 | The trap destination IP address is the IP address of the SNMP Server where the trap will be sent. Ensure that the SNMP IP is valid.   |
| <b>snmp_community_string</b><br>string Optional                 | An SNMP community string is a means of accessing statistics stored within a router or other device.<br><b>Default values:</b> <code>public</code>   |
| <b>ethernet 1/1/(1-63) config</b><br>string Required            | By default: <ul style="list-style-type: none"> <li>• Port description is provided.</li> <li>• Each interface is set to “up” state.</li> <li>• The fanout/breakout mode for 1/1/1 to 1/1/63 is as per the value set in the <code>breakout_value</code> variable.</li> <li>• Update the individual interfaces of the Dell PowerSwitch S5232F-ON.</li> <li>• The interfaces are from ethernet 1/1/1 to ethernet 1/1/63. By default, the breakout mode is set for 1/1/1 to 1/1/63.</li> <li>• Note: The playbooks will fail if any invalid configurations are entered.</li> </ul> |
| <b>save_changes_to_startup</b><br>boolean <sup>1</sup> Required | Change it to “true” only when you are certain that the updated configurations and commands are valid.<br><b>WARNING:</b> When set to “true”, the startup configuration file is updated. If incorrect configurations or commands are entered, the Ethernet switches may not operate as expected.<br>Choices: <ul style="list-style-type: none"> <li>• <code>false &lt;- Default</code></li> <li>• <code>true</code></li> </ul>   |

- When initializing a factory reset switch, the user needs to ensure DHCP is enabled and an IPv6 address is not assigned.
- The 65th port on a Z series switch cannot be split.

<sup>1</sup> Boolean parameters do not need to be passed with double or single quotes.

- Only odd ports support breakouts on Z9264F-ON. For more information, [click here](#).

---

**Note:** The `breakout_value` of a port can only be changed after un-splitting the port.

---

#### Running the playbook:

```
cd network

ansible-playbook ethernet_switch_config.yml -i inventory -e ethernet_switch_username="" -
↵e ethernet_switch_password=""
```

- Where `ethernet_switch_username` is the username used to authenticate into the switch.
- The inventory file should be a list of IPs separated by newlines. Check out the `switch_inventory` section in [Sample Files](#)
- Where `ethernet_switch_password` is the password used to authenticate into the switch.

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

## 2.8 Configuring storage

### Configuring Powervault storage

To configure powervault ME4 and ME5 storage arrays, follow the below steps:

Fill out all required parameters in `storage/powervault_input.yml`:

**Caution:** Do not remove or comment any lines in the `storage/powervault_input.yml` file.

| Parameter  | Details   |
|--|---|
| <b>powervault_protocol</b><br>string Required          | This variable indicates the network protocol used for data connectivity .<br><b>Default values:</b> sas   |
| <b>powervault_controller_mode</b><br>string Required   | This variable indicates the number of controllers available on the target powervault.<br>Choices: <ul style="list-style-type: none"> <li>• multi &lt;- default</li> <li>• single</li> </ul>   |
| <b>powervault_locale</b><br>string Optional            | Represents the selected language. Currently, only English is supported.<br><b>Default values:</b> English   |
| <b>powervault_system_name</b><br>string Optional       | The system name used to identify the PowerVault Storage device. The name should be less than 30 characters and must not contain spaces.<br><b>Default values:</b> Uninitialized_Name  |
| <b>powervault_snmp_notify_level</b><br>string Required | Select the SNMP notification levels for PowerVault Storage devices.<br><b>Default values:</b> none  |
| <b>powervault_pool_type</b><br>string Required         | This variable indicates the kind of pool created on the target powervault.<br>Choices: <ul style="list-style-type: none"> <li>• linear &lt;- default</li> <li>• virtual</li> </ul>  |
| <b>powervault_raid_levels</b><br>string Optional       | Enter the required RAID levels and the minimum / maximum number of disks for each RAID levels.<br>Choices: <ul style="list-style-type: none"> <li>• raid1 &lt;- default</li> <li>• raid5</li> <li>• raid6</li> <li>• raid10</li> </ul>  |
| <b>powervault_disk_range</b><br>string Required        | <b>Enter the range of disks in the format enclosure-number.disk-range,enclosure-number.disk-range. For example, to select disks 3 to 12 in enclosure 1 and to select disks 5 to 23 in enclosure 2, you must enter 1.3-12, 2.5-23.</b><br>A RAID 10 or 50 disk group with disks in sub-groups are separated by colons (with no spaces). RAID-10 example: 1.1-2:1.3-4:1.7,1.10 Note: Ensure that the entered disk location is empty and the Usage column lists the range as AVAIL. The disk range specified must be of the same vendor and they must have the same description.<br><b>Default values:</b> 0.0-1 |
| <b>powervault_disk_group_name</b><br>string Required   | Specifies the disk group name<br><b>Default values:</b> omnia   |
| <b>powervault_volumes</b><br>string Required           | Specify the volume details for powervault and NFS Server node. Multiple volumes can be defined as comma separated values. example: omnia homel omnia  |

Run the playbook:

```
cd storage
ansible-playbook powervault.yml -i inventory -e powervault_username="" -e powervault_
↪password=""
```

- Where the `inventory` refers to a list of all nodes separated by a newline.
- `powervault_username` and `powervault_password` are the credentials used to administrate the array.

**Note:** Once the storage is successfully set up, set `enable_omnia_nfs` (`input/omnia_config.yml`) to false and `omnia_usrhome_share` (`input/omnia_config.yml`) to an accessible share path in BeeGFS to use the path across the cluster for deployments.

### Configuring NFS servers

To configure an NFS server, enter the following parameters in `storage/nfs_server_input.yml`

| Parameter                                       | Details   |
|---|---|
| <b>powervault_ip</b><br>string Optional         | Mandatory field when <code>nfs</code> group is defined with an IP and <code>omnia</code> is required to configure <code>nfs</code> server. IP of Powervault connected to NFS Server should be provided. In a single run of <code>omnia</code> , only one NFS Server is configured. To configure multiple NFS Servers, add one IP in <code>nfs</code> group in a single run of <code>omnia.yml</code> and give variable values accordingly. To configure another <code>nfs</code> node, update variables and run <code>nfs_sas.yml</code>  |
| <b>powervault_volumes</b><br>JSON list Required | Specify the volume details for powervault and NFS Server node For multiple volumes, list of json with volume details should be provided. <ul style="list-style-type: none"> <li>• <b>server_share_path:</b> The path at which volume is mounted on <code>nfs</code> node</li> <li>• <b>server_export_options:</b> Default value is <code>rw,sync,no_root_squash</code> (unless specified otherwise). For a list of accepted options, <a href="#">click here</a></li> <li>• <b>client_shared_path:</b> The path at which volume is mounted on <code>manager</code>, <code>compute</code>, <code>login</code> node. This value is taken as <code>server_share_path</code> unless specified otherwise.</li> <li>• <b>client_mount_options:</b> Default value is <code>nosuid,rw,sync,hard,intr</code> (unless specified otherwise). For a list of accepted options, <a href="#">click here</a></li> </ul> <p>Must specify atleast 1 volume</p> <p><b>Default values:</b> `` - { name: omnia_home, server_share_path: /home/omnia_home, server_export_options: } ``</p> |

Run the playbook:

```
cd storage
ansible-playbook nfs_sas.yml -i /root/inventory -e powervault_username="xxxxxx" -e
↪powervault_password="xxxxxx"
```

- Where the inventory refers to a list of all nodes in the format of [NFS server inventory file](#)
- To set up NFS client services, [click here](#)

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

## 2.9 Running HPC benchmarks on omnia clusters

### 2.9.1 Install oneAPI for MPI jobs on Intel processors

This topic explains how to manually install oneAPI for MPI jobs. To install oneAPI automatically, [click here](#).

#### Pre-requisites

- An Omnia **slurm** cluster running with at least 2 nodes: 1 manager and 1 compute.
- Verify that the target nodes are in the booted state. For more information, [click here](#).

#### Download and install Intel oneAPI base toolkit & Intel oneAPI HPC toolkit to control plane

1. Create a DNF repository file in the /tmp directory as a normal user.

```
tee > /tmp/oneAPI.repo << EOF
[oneAPI]
name=Intel® oneAPI repository
baseurl=https://yum.repos.intel.com/oneapi
enabled=1
gpgcheck=1
repo_gpgcheck=1
gpgkey=https://yum.repos.intel.com/intel-gpg-keys/GPG-PUB-KEY-INTEL-SW-PRODUCTS.PUB
EOF
```

2. Move the newly created oneAPI.repo file to the YUM configuration directory.

```
sudo mv /tmp/oneAPI.repo /etc/yum.repos.d
```

3. Switch to the /install/post/otherpkgs/<Provision OS.Version>/x86\_64/custom\_software/Packages/ folder and execute:

For example: `cd /install/post/otherpkgs/rhels8.6.0/x86_64/custom_software/Packages/.`

```
dnf download intel-basekit --resolve --alldeps -y
dnf download intel-hpckit --resolve --alldeps -y
```

---

**Note:** Use `alldeps -y` to download **all** dependencies related to OneAPI.

---

4. Once downloaded, make sure there are  $\geq 270$  rpm packages in the /install/post/otherpkgs/rhels8.6.0/x86\_64/custom\_software/Packages/ directory.
5. Inside the /install/post/otherpkgs/<Provision OS.Version>/x86\_64/custom\_software folder, create a file named `update.otherpkgs.pkglist` with the contents:

```
custom_software/intel-basekit
custom_software/intel-hpckit
```



6. Go to `utils/os_package_update` and edit `package_update_config.yml`. For more information on the input parameters, [click here](#).
7. Run `package_update.yml` using: `ansible-playbook package_update.yml`
8. After execution is completed, verify that `intelhpckit` and `basekit` packages are on the nodes using: `rpm -qa | grep intel`

### To execute multi-node jobs

- Make sure to have NFS shares on each node.
- Copy slurm script to NFS share and execute it from there.
- Load all the necessary modules using module load:

```
module load mpi
module load pmi/pmix-x86_64
module load mkl
```

- If the commands/batch script are to be run over TCP instead of Infiniband ports, include the below line:

```
export FI_PROVIDER=tcp
```

Job execution can now be initiated.

---

**Note:** Ensure `runme_intel64_dynamic` is downloaded before running this command.

---

```
srunk -N 2 /mnt/nfs_shares/appshare/mkl/2023.0.0/benchmarks/mp_linpack/runme_intel64_
dynamic
```

For a batch job using the same parameters, the script would be:

```
#!/bin/bash
#SBATCH --job-name=testMPI
#SBATCH --output=output.txt
#SBATCH --partition=normal
#SBATCH --odelist=node00004.omnia.test,node00005.omnia.test

pwd; hostname; date
export FI_PROVIDER=tcp
module load pmi/pmix-x86_64
module use /opt/intel/oneapi/modulefiles
module load mkl
module load mpi

srunk /mnt/appshare/benchmarks/mp_linpack/runme_intel64_dynamic
date
```

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

## 2.9.2 Automate installation oneAPI on Intel processors for MPI jobs

This topic explains how to automatically update servers for MPI jobs. To manually install oneAPI, [click here](#).

### Pre-requisites

- `provision.yml` has been executed.
- An Omnia **slurm** cluster has been set up by `omnia.yml` running with at least 2 nodes: 1 manager and 1 compute.
- Verify that the target nodes are in the booted state. For more information, [click here](#).

### To run the playbook:

```
cd benchmarks
ansible-playbook intel_benchmark.yml -i inventory
```

### To execute multi-node jobs

- Make sure to have NFS shares on each node.
- Copy slurm script to NFS share and execute it from there.
- Load all the necessary modules using `module load`:

```
module load mpi
module load pmi/pmix-x86_64
module load mkl
```

- If the commands/batch script are to be run over TCP instead of Infiniband ports, include the below line:

```
export FI_PROVIDER=tcp
```

Job execution can now be initiated.

---

**Note:** Ensure `runme_intel64_dynamic` is downloaded before running this command.

---

```
srunk -N 2 /mnt/nfs_shares/appshare/mkl/2023.0.0/benchmarks/mp_linpack/runme_intel64_
↪dynamic
```

For a batch job using the same parameters, the script would be:

```
#!/bin/bash
#SBATCH --job-name=testMPI
#SBATCH --output=output.txt
#SBATCH --partition=normal
#SBATCH --odelist=node00004.omnia.test,node00005.omnia.test

pwd; hostname; date
export FI_PROVIDER=tcp
module load pmi/pmix-x86_64
module use /opt/intel/oneapi/modulefiles
module load mkl
module load mpi

srunk /mnt/appshare/benchmarks/mp_linpack/runme_intel64_dynamic
date
```

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

### 2.9.3 Open MPI AOCC HPL benchmark for AMD processors

This topic explains how to manually update servers for MPI jobs. To automatically install pmix and configure slurm, [click here](#).

#### Prerequisites

1. Provision the cluster and install slurm on all cluster nodes.
2. Dependent packages have to be installed on the cluster nodes using the following steps.:
  - i. Download the dependent packages on the control plane.:

- a. Create a package list.:

```
cd /install/post/otherpkgs/<os_version>/x86_64/custom_software/
cat openmpi.pkglist
```

- b. Enter the following contents into openmpi.pkglist:

```
custom_software/pmix-devel
custom_software/libevent-devel
```

- c. Download the packages:

```
cd packages
dnf download pmix-devel --resolve --alldeps
dnf download libevent-devel --resolve --alldeps
```

- ii. Push the packages to the cluster nodes:

- a. Update the package\_list variable in the utils/os\_package\_update/package\_update\_config.yml file and save it.

```
package_list: "/install/post/otherpkgs/<os_version>/x86_64/custom_software/
↳openmpi.pkglist"
```

- b. Update the cluster nodes by running the package\_update.yml playbook.

```
ansible-playbook package_update.yml
```

3. OpenMPI and aocc-compiler-\*.tar should be installed and compiled with slurm on all cluster nodes or should be available on the NFS share.

#### Note:

- Omnia currently supports pmix version2, pmix\_v2.
- While compiling OpenMPI, include pmix, slurm, hwloc and, libevent as shown in the below sample command:

```
./configure --prefix=/home/omnia-share/openmpi-4.1.5 --enable-mpi1-compatibility --
↳enable-orterun-prefix-by-default --with-slurm=/usr --with-pmix=/usr --with-
↳libevent=/usr --with-hwloc=/usr --with-ucx CC=clang CXX=clang++ FC=flang 2>&1 |
↳tee config.out
```

**To execute multi-node jobs**

1. Update the following parameters in `/etc/slurm/slurm.conf`:

```
SelectType=select/cons_tres
SelectTypeParameters=CR_Core
TaskPlugin=task/affinity,task/cgroup
```

2. Restart `slurmd.service` on all compute nodes.

```
systemctl stop slurmd
systemctl start slurmd
```

3. Once the service restarts on the compute nodes, restart `slurmctld.service` on the manager node.

```
systemctl stop slurmctld.service
systemctl start slurmctld.service
```

4. Job execution can now be initiated.

For a job to run on multiple nodes (10.5.0.4 and 10.5.0.5) where OpenMPI is compiled and installed on the NFS share (`/home/omnia-share/openmpi/bin/mpirun`), the job can be initiated as below: .. note:: Ensure `amd-zen-hpl-2023_07_18` is downloaded before running this command.

```
srun -N 2 --mpi=pmix_v2 -n 2 ./amd-zen-hpl-2023_07_18/xhpl
```

For a batch job using the same parameters, the script would be:

```
#!/bin/bash

#SBATCH --job-name=test

#SBATCH --output=test.log

#SBATCH --partition=normal

#SBATCH -N 3

#SBATCH --time=10:00

#SBATCH --ntasks=2


source /home/omnia-share/setenv_AOCC.sh

export PATH=$PATH:/home/omnia-share/openmpi/bin

export LD_LIBRARY_PATH=$LD_LIBRARY_PATH:/home/omnia-share/openmpi/lib

srun --mpi=pmix_v2 ./amd-zen-hpl-2023_07_18/xhpl
```

Alternatively, to use `mpirun`, the script would be:

```
#!/bin/bash

#SBATCH --job-name=test

#SBATCH --output=test.log

#SBATCH --partition=normal

#SBATCH -N 3

#SBATCH --time=10:00

#SBATCH --ntasks=2


source /home/omnia-share/setenv_AOCC.sh

export PATH=$PATH:/home/omnia-share/openmpi/bin

export LD_LIBRARY_PATH=$LD_LIBRARY_PATH:/home/omnia-share/openmpi/lib

/home/omnia-share/openmpi/bin/mpirun --map-by ppr:1:node -np 2 --display-map --
↪oversubscribe --mca orte_keep_fqdn_hostnames 1 ./xhpl
```

**Note:** The above scripts are samples that can be modified as required. Ensure that `--mca orte_keep_fqdn_hostnames 1` is included in the `mpirun` command in `sbatch` scripts. Omnia maintains all hostnames in FQDN format. Failing to include `--mca orte_keep_fqdn_hostnames 1` may cause job initiation to fail.

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

## 2.9.4 Installing pmix and updating slurm configuration for AMD processors

This topic explains how to automatically update AMD servers for MPI jobs. To manually install `pmix` and update the `slurm` configuration, [click here](#).

### Pre-requisites

- `provision.yml` has been executed.
- An Omnia **slurm** cluster has been set up by `omnia.yml` running with at least 2 nodes: 1 manager and 1 compute.
- Verify that the target nodes are in the booted state. For more information, [click here](#).

### To run the playbook:

```
cd benchmarks
ansible-playbook amd_benchmark.yml -i inventory
```

### To execute multi-node jobs

- OpenMPI and aocc-compiler-\*.tar should be installed and compiled with slurm on all cluster nodes or should be available on the NFS share.

---

**Note:**

- Omnia currently supports pmix version2, pmix\_v2.
- While compiling OpenMPI, include pmix, slurm, hwloc and, libevent as shown in the below sample command:

```
./configure --prefix=/home/omnia-share/openmpi-4.1.5 --enable-mpi1-compatibility --  
↪ enable-orterun-prefix-by-default --with-slurm=/usr --with-pmix=/usr --with-  
↪ libevent=/usr --with-hwloc=/usr --with-ucx CC=clang CXX=clang++ FC=flang 2>&1 |  
↪ tee config.out
```

- For a job to run on multiple nodes (10.5.0.4 and 10.5.0.5) where OpenMPI is compiled and installed on the NFS share (/home/omnia-share/openmpi/bin/mpirun), the job can be initiated as below:

---

**Note:** Ensure amd-zen-hpl-2023\_07\_18 is downloaded before running this command.

---

```
srun -N 2 --mpi=pmix_v2 -n 2 ./amd-zen-hpl-2023_07_18/xhpl
```

For a batch job using the same parameters, the script would be:

```
#!/bin/bash  
  
#SBATCH --job-name=test  
  
#SBATCH --output=test.log  
  
#SBATCH --partition=normal  
  
#SBATCH -N 3  
  
#SBATCH --time=10:00  
  
#SBATCH --ntasks=2  
  
  
  
source /home/omnia-share/setenv_AOCC.sh  
  
export PATH=$PATH:/home/omnia-share/openmpi/bin  
  
export LD_LIBRARY_PATH=$LD_LIBRARY_PATH:/home/omnia-share/openmpi/lib  
  
srun --mpi=pmix_v2 ./amd-zen-hpl-2023_07_18/xhpl
```

Alternatively, to use mpirun, the script would be:

```
#!/bin/bash

#SBATCH --job-name=test

#SBATCH --output=test.log

#SBATCH --partition=normal

#SBATCH -N 3

#SBATCH --time=10:00

#SBATCH --ntasks=2


source /home/omnia-share/setenv_AOCC.sh

export PATH=$PATH:/home/omnia-share/openmpi/bin

export LD_LIBRARY_PATH=$LD_LIBRARY_PATH:/home/omnia-share/openmpi/lib

/home/omnia-share/openmpi/bin/mpirun --map-by ppr:1:node -np 2 --display-map --
↪oversubscribe --mca orte_keep_fqdn_hostnames 1 ./xhpl
```

**Note:** The above scripts are samples that can be modified as required. Ensure that `--mca orte_keep_fqdn_hostnames 1` is included in the `mpirun` command in `sbatch` scripts. Omnia maintains all hostnames in FQDN format. Failing to include `--mca orte_keep_fqdn_hostnames 1` may cause job initiation to fail.

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

## 2.9.5 Containerized HPC benchmark execution

Use this playbook to download docker images and pull images onto cluster nodes using `apptainer`.

1. Ensure that the cluster has been [provisioned by the provision tool](#). and the cluster has been set up using `omnia.yml`.
2. Enter the following variables in `utils/hpc_apptainer_job_execution/hpc_apptainer_job_execution_config.yml`:

| Parameter                                       | Details   |
|---|---|
| <b>hpc_apptainer_image</b> JSON list Required   | <ul style="list-style-type: none"> <li>Docker image details to be downloaded in to cluster nodes using apptainer to create a sif file.</li> <li>Example (for single image): <pre>hpc_apptainer_image: - { image_url: "docker.io/intel/ ↪oneapi-hpckit:latest" }</pre> </li> <li>Example (for multiple images): <pre>hpc_apptainer_image: - { image_url: "docker.io/intel/ ↪oneapi-hpckit:latest" }  - { image_url: "docker.io/ ↪tensorflow/tensorflow:latest" }</pre> </li> <li>If provided, docker credentials in <code>omnia_config.yml</code>, it will be used for downloading docker images.</li> </ul> |
| <b>hpc_apptainer_path</b><br>string<br>Required | <ul style="list-style-type: none"> <li>Directory to filepath for storing apptainer sif files on cluster nodes.</li> <li>It is recommended to use a directory inside a shared path that is accessible to all cluster nodes.</li> <li><b>Default value:</b> <code>"/home/omnia-share/softwares/apptainer"</code></li> </ul>   |

To run the playbook:

```
cd utils/hpc_apptainer_job_execution
ansible-playbook hpc_apptainer_job_execution.yml -i inventory
```

**Note:** Use the inventory file format specified under [Sample Files](#).

HPC apptainer jobs can be initiated on a slurm cluster using the following sample command:

```
srun -N 3 --mpi=pmi2 --ntasks=4 apptainer run /home/omnia-share/softwares/apptainer/
↪oneapi-hpckit_latest.sif hostname
```

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).



## 2.10 Uninstalling the provision tool

Use this script to undo all the changes made by the provision tool. For a list of actions taken by the provision tool, [click here](#).

To run the script:

```
cd utils
ansible-playbook control_plane_cleanup.yml
```

**Caution:**

- When re-provisioning your cluster (that is, re-running the `provision.yml` playbook) after a clean-up, ensure to use a different `admin_nic_subnet` in `input/provision_config.yml` to avoid a conflict with newly assigned servers. Alternatively, disable any OS available in the `Boot Option Enable/Disable` section of your BIOS settings (BIOS Settings > Boot Settings > UEFI Boot Settings) on all target nodes.
- On subsequent runs of `provision.yml`, if users are unable to log into the server, refresh the ssh key manually and retry.

```
ssh-keygen -R <node IP>
```

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).



## FEATURES

From Omnia 1.4, all of Omnia's many features are available individually. Specific playbooks allow users to choose different features and customize their deployment journey specifically to their needs.

Below is a list of all Omnia's features:

### 3.1 Network

In your HPC cluster, connect the Mellanox InfiniBand switches using the Fat-Tree topology. In the fat-tree topology, switches in layer 1 are connected through the switches in the upper layer, i.e., layer 2. And, all the cluster nodes in the cluster, such as PowerEdge servers and PowerVault storage devices, are connected to switches in layer 1. With this topology in place, we ensure that a 1x1 communication path is established between the cluster nodes. For more information on the fat-tree topology, see [Designing an HPC cluster with Mellanox infiniband-solutions](#).

---

**Note:**

- From Omnia 1.4, the Subnet Manager runs on the target Infiniband switches and not the control plane.
  - When `ib_nic_subnet` is provided in `input/provision_config.yml`, the infiniband NIC on target nodes are assigned IPv4 addresses within the subnet without user intervention during the execution of `provision.yml`.
- 

Some of the network features Omnia offers are:

1. Mellanox OFED
2. Infiniband switch configuration

To install OFED drivers, enter all required parameters in `input/network_config.yml`:

| Name  | Description   |
|---|---|
| <b>mlnx_ofed_offline_path</b><br>[optional] string        | Absolute path to local copy of .tgz file containing mlnx_ofed package. The package can be downloaded from <a href="https://network.nvidia.com/products/infiniband-drivers/linux/mlnx_ofed/">https://network.nvidia.com/products/infiniband-drivers/linux/mlnx_ofed/</a> . |
| <b>mlnx_ofed_version</b><br>[optional] string             | Indicates the version of mlnx_ofed to be downloaded. If mlnx_ofed_offline_path is not given, declaring this variable is mandatory.<br><b>Default value:</b> 5.7-1.0.2.0   |
| <b>mlnx_ofed_add_kernel_support</b><br>[required] boolean | Indicates whether the kernel needs to be upgraded to be compatible with mlnx_ofed.<br><b>Choices:</b> <ul style="list-style-type: none"><li>• false &lt;- Default</li><li>• true</li></ul>  |

To run the script:

```
cd network
ansible-playbook network.yml
```

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

## 3.2 Centralized authentication on the cluster

The security feature allows users to set up FreeIPA and LDAP to help authenticate into HPC clusters.

---

**Note:**

- Nodes provisioned using the Omnia provision tool do not require a RedHat subscription to run `security.yml` on RHEL target nodes.
  - For RHEL target nodes not provisioned by Omnia, ensure that RedHat subscription is enabled on all target nodes. Every target node will require a RedHat subscription.
-

### 3.2.1 Configuring FreeIPA/LDAP security

Enter the following parameters in `input/security_config.yml`.

| Parameter  | Details  |
|--|--|
| <b>freeipa_required</b><br>boolean Optional          | Boolean indicating whether FreeIPA is required or not.<br>Choices: <ul style="list-style-type: none"> <li>• <code>true</code> &lt;- Default</li> <li>• <code>false</code></li> </ul>   |
| <b>realm_name</b><br>string Optional                 | Sets the intended realm name.<br><b>Default values:</b> <code>OMNIA.TEST</code>  |
| <b>directory_manager_password</b><br>string Optional | Password authenticating admin level access to the Directory for system management tasks. It will be added to the instance of directory server created for IPA. Required Length: 8 characters. The password must not contain -, , ' , " ,   |
| <b>kerberos_admin_password</b><br>string Optional    | "admin" user password for the IPA server on RockyOS.   |
| <b>ldap_required</b><br>boolean Optional             | Boolean indicating whether ldap client is required or not.<br>Choices: <ul style="list-style-type: none"> <li>• <code>false</code> &lt;- Default</li> <li>• <code>true</code></li> </ul>   |
| <b>domain_name</b><br>string Optional                | Sets the intended domain name.<br><b>Default values:</b> <code>omnia.test</code>   |
| <b>ldap_server_ip</b><br>string Optional             | LDAP server IP. Required if <code>ldap_required</code> is true. There should be an explicit LDAP server running on this IP.  |
| <b>ldap_connection_type</b><br>string Optional       | For a TLS connection, provide a valid certification path. For an SSL connection, ensure port 636 is open.<br><b>Default values:</b> <code>TLS</code>   |
| <b>ldap_ca_cert_path</b><br>string Optional          | This variable accepts Server Certificate Path. Make sure certificate is present in the path provided. The certificate should have .pem or .crt extension. This variable is mandatory if connection type is TLS.<br><b>Default values:</b> <code>/etc/openldap/certs/omnialdap.pem</code>           |
| <b>user_home_dir</b><br>string Optional              | This variable accepts the user home directory path for ldap configuration. If nfs mount is created for user home, make sure you provide the LDAP users mount home directory path.<br><b>Default values:</b> <code>/home</code>   |
| <b>ldap_bind_username</b><br>string Optional         | If LDAP server is configured with bind dn then bind dn user to be provided. If this value is not provided (when bind is configured in server) then ldap authentication fails. Omnia does not validate this input. Ensure that it is valid and proper.<br><b>Default values:</b> <code>admin</code> |
| <b>ldap_bind_password</b><br>string Optional         | If LDAP server is configured with bind dn then bind dn password to be provided. If this value is not provided (when bind is configured in server) then ldap authentication fails. Omnia does not validate this input. Ensure   |

---

**Note:** When `ldap_required` is true, `freeipa_required` has to be false. Conversely, when `freeipa_required` is true, `ldap_required` has to be false.

---

### 3.2.2 Configuring login node security

Enter the following parameters in `input/login_node_security_config.yml`.

| Variable  | Details  |
|---|--|
| <b>max_failures</b><br>integer Optional             | The number of login failures that can take place before the account is locked out.<br><b>Default values:</b> 3   |
| <b>failure_reset_interval</b><br>integer Optional   | Period (in seconds) after which the number of failed login attempts is reset. Min value: 30; Max value: 60.<br><b>Default values:</b> 60   |
| <b>lockout_duration</b><br>integer Optional         | Period (in seconds) for which users are locked out. Min value: 5; Max value: 10.<br><b>Default values:</b> 10  |
| <b>session_timeout</b><br>integer Optional          | User sessions that have been idle for a specific period can be ended automatically. Min value: 90; Max value: 180.<br><b>Default values:</b> 180   |
| <b>alert_email_address</b><br>string Optional       | Email address used for sending alerts in case of authentication failure. When blank, authentication failure alerts are disabled. Currently, only one email ID is accepted.   |
| <b>user</b><br>string Optional                      | Access control list of users. Accepted formats are <code>username@ip</code> ( <code>root@1.2.3.4</code> ) or username ( <code>root</code> ). Multiple users can be separated using whitespaces.  |
| <b>allow_deny</b><br>string Optional                | This variable decides whether users are to be allowed or denied access. Ensure that AllowUsers or DenyUsers entries on sshd configuration file are not commented.<br>Choices: <ul style="list-style-type: none"> <li>• <code>allow</code> &lt;- Default</li> <li>• <code>deny</code></li> </ul>                |
| <b>restrict_program_support</b><br>boolean Optional | This variable is used to disable services. Root access is mandatory.<br>Choices: <ul style="list-style-type: none"> <li>• <code>false</code> &lt;- Default</li> <li>• <code>true</code></li> </ul>   |
| <b>restrict_softwares</b><br>string Optional        | List of services to be disabled (Comma-separated). Example: <code>'telnet,lpd,bluetooth'</code><br>Choices: <ul style="list-style-type: none"> <li>• <code>telnet</code></li> <li>• <code>lpd</code></li> <li>• <code>bluetooth</code></li> <li>• <code>rlogin</code></li> <li>• <code>rexec</code></li> </ul> |



### 3.2.3 Installing LDAP Client

Manager and compute nodes will have LDAP client installed and configured if `ldap_required` is set to true. The login node does not have LDAP client installed.

**Caution:** No users/groups will be created by Omnia.

#### FreeIPA installation on the NFS node

IPA services are used to provide account management and centralized authentication.

To customize your installation of FreeIPA, enter the following parameters in `input/security_config.yml`.

| Input Parameter                  | Definition  | Variable value   |
|----------------------------------|---|--|
| <code>kerberos_admin</code>      | “admin” user password for the IPA server on RockyOS and RedHat. | The password can be found in the file <code>input/security_config.yml</code> .   |
| <code>ipa_server_hostname</code> | The hostname of the IPA server                                  | The hostname can be found on the manager node.   |
| <code>domain_name</code>         | Domain name   | The domain name can be found in the file <code>input/security_config.yml</code> .  |
| <code>ipa_server_ip</code>       | The IP address of the IPA server                                | The IP address can be found on the IPA server on the manager node using the <code>ip a</code> command. This IP address should be accessible from the NFS node. |

To set up IPA services for the NFS node in the target cluster, run the following command from the `utils/cluster` folder on the control plane:

```
cd utils/cluster
ansible-playbook install_ipa_client.yml -i inventory -e kerberos_admin_password="" -e ipa_server_hostname="" -e domain_name="" -e ipa_server_ipaddress=""
```

#### Hostname requirements

- The Hostname should not contain the following characters: , (comma), . (period) or \_ (underscore). However, the **domain name** is allowed commas and periods.
- The Hostname cannot start or end with a hyphen (-).
- No upper case characters are allowed in the hostname.
- The hostname cannot start with a number.
- The hostname and the domain name (that is: `hostname000000x.domain.xxx`) cumulatively cannot exceed 64 characters. For example, if the `node_name` provided in `input/provision_config.yml` is ‘node’, and the `domain_name` provided is ‘omnia.test’, Omnia will set the hostname of a target cluster node to ‘node00001.omnia.test’. Omnia appends 6 digits to the hostname to individually name each target node.

**Note:** Use the format specified under [NFS inventory in the Sample Files](#) for inventory.

## Running the security role

Run:

```
cd security
ansible-playbook security.yml -i inventory
```

The inventory should contain compute, manager, login as per the inventory file in [samplefiles](#). The inventory file is case-sensitive. Follow the casing provided in the sample file link.

- To enable security features on the login node, ensure that `enable_secure_login_node` in `input/security_config.yml` is set to true.
- To customize the security features on the login node, fill out the parameters in `input/login_node_security_config.yml`.
- If a subsequent run of `security.yml` fails, the `security_config.yml` file will be unencrypted.

**Caution:** No users/groups will be created by Omnia.

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

## 3.3 Shared and distributed storage deployment

The storage role allows users to configure PowerVault Storage devices, BeeGFS and NFS services on the cluster.

First, enter all required parameters in `input/storage_config.yml`

| Name   | Details   |
|--|---|
| <b>nfs_client_params</b><br>JSON list Optional | <p>If NFS client services are to be deployed, enter the configuration required here in JSON format. The server_ip provided should have an explicit NFS server running. If left blank, no NFS configuration takes place. Possible values include:</p> <ol style="list-style-type: none"> <li>1. Single NFS file system: A single filesystem from a single NFS server is mounted.<br/>           Sample value: <code>- { server_ip: xx.xx.xx.xx, server_share_path: "/mnt/share", client_share_path: "/mnt/client", client_mount_options: "nosuid,rw,sync,hard,intr" }</code></li> <li>2. Multiple Mount NFS file system: Multiple filesystems from a single NFS server are mounted. Sample values:               <ul style="list-style-type: none"> <li>- <code>{ server_ip: xx.xx.xx.xx, server_share_path: "/mnt/server1", client_share_path: "/mnt/client1", client_mount_options: "nosuid,rw,sync,hard,intr" }</code></li> <li>- <code>{ server_ip: xx.xx.xx.xx, server_share_path: "/mnt/server2", client_share_path: "/mnt/client2", client_mount_options: "nosuid,rw,sync,hard,intr" }</code></li> </ul> </li> <li>3. Multiple NFS file systems: Multiple filesystems are mounted from multiple servers. Sample Values:               <ul style="list-style-type: none"> <li>- <code>{ server_ip: zz.zz.zz.zz, server_share_path: "/mnt/share1", client_share_path: "/mnt/client1", client_mount_options: "nosuid,rw,sync,hard,intr" }</code></li> <li>- <code>{ server_ip: xx.xx.xx.xx, server_share_path: "/mnt/share2", client_share_path: "/mnt/client2", client_mount_options: "nosuid,rw,sync,hard,intr" }</code></li> <li>- <code>{ server_ip: yy.yy.yy.yy, server_share_path: "/mnt/share3", client_share_path: "/mnt/client3", client_mount_options: "nosuid,rw,sync,hard,intr" }</code></li> </ul> </li> </ol> <p><b>Default value:</b> <code>{ server_ip: , server_share_path: , client_share_path: , client_mount_options: }</code></p> |
| <b>beegfs_support</b><br>boolean Optional      | <p>This variable is used to install beegfs-client on compute and manager nodes</p> <p>Choices:</p> <ul style="list-style-type: none"> <li>• false &lt;- Default</li> <li>• true</li> </ul>  |
| <b>beegfs_rdma_support</b><br>boolean Optional | <p>This variable is used if user has RDMA-capable network hardware (e.g., InfiniBand)</p> <p>Choices:</p> <ul style="list-style-type: none"> <li>• false &lt;- Default</li> <li>• true</li> </ul>   |

**Note:** If `storage.yml` is run with the `input/storage_config.yml` filled out, BeeGFS and NFS client will be set up.

---

### Installing BeeGFS Client

- If the user intends to use BeeGFS, ensure that a BeeGFS cluster has been set up with `beegfs-mgmt`, `beegfs-meta`, `beegfs-storage` services running.

Ensure that the following ports are open for TCP and UDP connectivity:

| Port | Service   |
|------|---|
| 8008 | Management service ( <code>beegfs-mgmt</code> ) |
| 8003 | Storage service ( <code>beegfs-storage</code> ) |
| 8004 | Client service ( <code>beegfs-client</code> )   |
| 8005 | Metadata service ( <code>beegfs-meta</code> )   |
| 8006 | Helper service ( <code>beegfs-helper</code> )   |

To open the ports required, use the following steps:

1. `firewall-cmd --permanent --zone=public --add-port=<port number>/tcp`
  2. `firewall-cmd --permanent --zone=public --add-port=<port number>/udp`
  3. `firewall-cmd --reload`
  4. `systemctl status firewalld`
- Ensure that the nodes in the inventory have been assigned **only** these roles: `manager` and `compute`.

---

#### Note:

- When working with RHEL, ensure that the BeeGFS configuration is supported using the [link here](#).
- If the BeeGFS server (MGMTD, Meta, or storage) is running BeeGFS version 7.3.1 or higher, the security feature on the server should be disabled. Change the value of `connDisableAuthentication` to `true` in `/etc/beegfs/beegfs-mgmt.conf`, `/etc/beegfs/beegfs-meta.conf` and `/etc/beegfs/beegfs-storage.conf`. Restart the services to complete the task:

```
systemctl restart beegfs-mgmt
systemctl restart beegfs-meta
systemctl restart beegfs-storage
systemctl status beegfs-mgmt
systemctl status beegfs-meta
systemctl status beegfs-storage
```

---

### NFS bolt-on

- Ensure that an external NFS server is running. NFS clients are mounted using the external NFS server's IP.
- Fill out the `nfs_client_params` variable in the `storage_config.yml` file in JSON format using the samples provided above.
- This role runs on `manager`, `compute` and `login` nodes.
- Make sure that `/etc/exports` on the NFS server is populated with the same paths listed as `server_share_path` in the `nfs_client_params` in `omnia_config.yml`.

- Post configuration, enable the following services (using this command: `firewall-cmd --permanent --add-service=<service name>`) and then reload the firewall (using this command: `firewall-cmd --reload`).
  - nfs
  - rpc-bind
  - mountd
- Omnia supports all NFS mount options. Without user input, the default mount options are `no-suid,rw,sync,hard,intr`. For a list of mount options, [click here](#).
- The fields listed in `nfs_client_params` are:
  - `server_ip`: IP of NFS server
  - `server_share_path`: Folder on which NFS server mounted
  - `client_share_path`: Target directory for the NFS mount on the client. If left empty, respective `server_share_path` value will be taken for `client_share_path`.
  - `client_mount_options`: The mount options when mounting the NFS export on the client. Default value: `nosuid,rw,sync,hard,intr`.
- There are 3 ways to configure the feature:
  1. **Single NFS node** : A single NFS filesystem is mounted from a single NFS server. The value of `nfs_client_params` would be:

```
- { server_ip: xx.xx.xx.xx, server_share_path: "/mnt/share", client_share_path:
  ↪"/mnt/client", client_mount_options: "nosuid,rw,sync,hard,intr" }
```

2. **Multiple Mount NFS Filesystem**: Multiple filesystems are mounted from a single NFS server. The value of `nfs_client_params` would be:

```
- { server_ip: xx.xx.xx.xx, server_share_path: "/mnt/server1", client_share_
  ↪path: "/mnt/client1", client_mount_options: "nosuid,rw,sync,hard,intr" }
- { server_ip: xx.xx.xx.xx, server_share_path: "/mnt/server2", client_share_
  ↪path: "/mnt/client2", client_mount_options: "nosuid,rw,sync,hard,intr" }
```

3. **Multiple NFS Filesystems**: Multiple filesystems are mounted from multiple NFS servers. The value of `nfs_client_params` would be:

```
- { server_ip: xx.xx.xx.xx, server_share_path: "/mnt/server1", client_share_
  ↪path: "/mnt/client1", client_mount_options: "nosuid,rw,sync,hard,intr" }
- { server_ip: yy.yy.yy.yy, server_share_path: "/mnt/server2", client_share_
  ↪path: "/mnt/client2", client_mount_options: "nosuid,rw,sync,hard,intr" }
- { server_ip: zz.zz.zz.zz, server_share_path: "/mnt/server3", client_share_
  ↪path: "/mnt/client3", client_mount_options: "nosuid,rw,sync,hard,intr" }
```

#### To run the playbook:

```
cd omnia/storage
ansible-playbook storage.yml -i inventory
```

(Where inventory refers to the [inventory file](#) listing manager, login\_node and compute nodes.)

**Note:** If a subsequent run of `storage.yml` fails, the `storage_config.yml` file will be unencrypted.

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

## 3.4 GPU accelerator configuration

The accelerator role allows users to set up the [AMD ROCm](#) platform or the [CUDA Nvidia toolkit](#). These tools allow users to unlock the potential of installed GPUs.

Enter all required parameters in `input/accelerator_config.yml`.

---

**Note:**

- Nodes provisioned using the Omnia provision tool do not require a RedHat subscription to run `accelerator.yml` on RHEL target nodes.
  - For RHEL target nodes not provisioned by Omnia, ensure that RedHat subscription is enabled on all target nodes. Every target node will require a RedHat subscription.
  - If `cuda_toolkit_path` is provided in `input/provision_config.yml` and NVIDIA GPUs are available on the target nodes, CUDA packages will be deployed post provisioning without user intervention during the execution of `provision.yml`.
  - AMD ROCm driver installation is not supported by Omnia on Rocky cluster nodes.
- 

To install all the latest GPU drivers and toolkits, run:

```
cd accelerator
ansible-playbook accelerator.yml -i inventory
```

(where inventory consists of manager, cluster and login nodes)

**The following configurations take place when running `accelerator.yml`**

- i. Servers with AMD GPUs are identified and the latest GPU drivers and ROCm platforms are downloaded and installed.
- ii. Servers with NVIDIA GPUs are identified and the specified CUDA toolkit is downloaded and installed.
- iii. For the rare servers with both NVIDIA and AMD GPUs installed, all the above mentioned download-ables are installed to the server.
- iv. Servers with neither GPU are skipped.

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

## 3.5 Offline repositories for the cluster

- The airgap feature will help create offline repositories on control plane which all the cluster nodes will access. This will remove the overhead of subscribing all the cluster nodes to RHEL.
- Currently, `airgap.yml` only updates RHEL repositories.

`airgap.yml` runs based on the following parameters in `input/provision_config.yml`:

| Parameter  | Details   |
|--|---|
| <b>update_repos</b><br>boolean<br>Required                   | <ul style="list-style-type: none"> <li>Indicates whether <code>provision.yml</code> will update offline RHEL repos (applicable from the second run of <code>provision.yml</code>)</li> <li>In the first execution of <code>provision.yml</code>, Omnia updates the BaseOS, Appstream and CRB repos.</li> <li>If <code>update_repos: false</code>, none of the repos required for cluster nodes will be updated provided the repos are already available.</li> <li>If <code>update_repos: true</code>, BaseOS, Appstream and CRB repos created for cluster nodes will be updated</li> </ul> <p>Choices:</p> <ul style="list-style-type: none"> <li><code>false</code> &lt;- Default</li> <li><code>true</code></li> </ul>  |
| <b>rhel_repo_alphabetical_folders</b><br>boolean<br>Required | <ul style="list-style-type: none"> <li>Indicates whether the packages in the local repos or subscription repos are ordered in alphabetical directories.</li> <li>For RHEL 8, when subscription is activated, this variable should be set to true.</li> </ul> <p>Choices:</p> <ul style="list-style-type: none"> <li><code>false</code> &lt;- Default</li> <li><code>true</code></li> </ul>  |
| <b>rhel_repo_local_path</b><br>JSON list<br>Optional         | <ul style="list-style-type: none"> <li>The repo path and names of the software repository to be configured on the cluster nodes.</li> <li>Provide the repo data file path, which ends with <code>.repo</code> extension in <code>repo_url</code> parameter.</li> <li>Provide a <b>valid</b> url for BaseOS and AppStream repositories.</li> <li>This variable should be filled if control plane OS is RHEL and subscription is not activated.</li> <li>This variable should be filled if the control plane OS is Rocky and the <code>provision_os</code> is <code>rhel</code>.</li> </ul> <p><b>Note:</b> Omnia does not validate the <code>repo_url</code> provided. Invalid entries will cause <code>provision.yml</code> to fail.</p> <p>Sample value:</p> <pre> - { repo: "AppStream", repo_url: "http://xx.yy.zz/pub/Distros/RedHat/RHEL8/8.8/RHEL-8-appstream.repo", repo_name: "RHEL-8-appstream-partners" }  - { repo: "BaseOS", repo_url: "http://xx.yy.zz/pub/Distros/RedHat/RHEL8/8.8/RHEL-8-baseos.repo", repo_name: "RHEL-8-baseos-partners" }</pre> |

airgap.yml is internally called when provision.yml is executed. Alternatively, run the following commands:

```
cd airgap
ansible-playbook airgap.yml
```

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

## 3.6 Additional utilities

The Utilities role allows users to set up certain tasks such as

### 3.6.1 Extra Packages for Enterprise Linux (EPEL)

This script is used to install the following packages:

1. PDSH
2. PDSH RCMD SSH
3. clustershell

To run the script:

```
cd omnia/utils
ansible-playbook install_hpc_thirdparty_packages.yml -i inventory
```

Where the inventory refers to a file listing all manager and compute nodes per the format provided in [inventory file](#).

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

### 3.6.2 Updating kernels on RHEL (with subscription)

#### Pre-requisites

1. Subscription should be available on nodes.
2. Kernels to be upgraded should be available. To verify the status of the kernels, use `yum list kernel`.
3. The input kernel revision cannot be a RHEL 7.x supported kernel version. e.g. “3.10.0-54.0.1” to “3.10.0-1160”.
4. Input needs to be passed during execution of the playbook.

#### Executing the Kernel Upgrade:

Via CLI:

```
cd omnia/utils
ansible-playbook kernel_upgrade.yml -i inventory -e rhsm_kernel_version=x.xx.x-xxxx
```

Where the inventory refers to a file listing all manager and compute nodes per the format provided in [inventory file](#). The inventory file is case-sensitive. Follow the casing provided in the sample file link.

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).



### 3.6.3 Red Hat Subscription

#### Required Parameters

| Name  | Description   |
|---|---|
| <b>redhat_subscription_method</b><br>string Optional          | Method to use for activation of subscription management. If Satellite, the role will determine the Satellite Server version (5 or 6) and take the appropriate registration actions.<br>Choices:<br><ul style="list-style-type: none"> <li>• portal &lt;- Default</li> <li>• satellite</li> </ul>  |
| <b>redhat_subscription_release</b><br>string Optional         | RHEL release version (e.g. 8.1)   |
| <b>redhat_subscription_force_register</b><br>boolean Optional | Register the system even if it is already registered.<br>Choices:<br><ul style="list-style-type: none"> <li>• false &lt;- Default</li> <li>• true</li> </ul>  |
| <b>redhat_subscription_pool_ids</b><br>string Optional        | <b>Specify subscription pool IDs to consume.</b> A pool ID may be specified as a string - just the pool ID (ex. 0123456789abcdef0123456789abcdef) or as a dict with the pool ID as the key, and a quantity as the value<br>If the quantity is provided, it is used to consume multiple entitlements from a pool (the pool must support this). |
| <b>redhat_subscription_repos</b><br>string Optional           | The list of repositories to enable or disable. When providing multiple values, a YAML list or a comma separated list are accepted.  |
| <b>redhat_subscription_repos_state</b><br>string Optional     | The state of all repos in redhat_subscription_repos.<br>Choices:<br><ul style="list-style-type: none"> <li>• enabled &lt;- Default</li> <li>• disabled</li> </ul>   |
| <b>redhat_subscription_repos_purge</b><br>boolean Optional    | This parameter disables all currently enabled repositories that are not specified in redhat_subscription_repos. Only set this to true if the redhat_subscription_repos field has multiple repos.<br>Choices:<br><ul style="list-style-type: none"> <li>• false &lt;- Default</li> <li>• true</li> </ul>                                       |
| <b>redhat_subscription_server_hostname</b><br>string Optional | FQDN of subscription server. Mandatory field if redhat_subscription_method is set to satellite.<br><b>Default values:</b> subscription.rhn.redhat.com   |
| <b>redhat_subscription_port</b><br>integer Optional           | Port to use when connecting to subscription server. Set 443 for Satellite or RHN. If capsule is used, set 8443.<br>Choices:<br><ul style="list-style-type: none"> <li>• 443 &lt;- Default</li> <li>• 8443</li> </ul>  |
| <b>redhat_subscription_insecure</b><br>boolean Optional       | Disable certificate validation.<br>Choices:<br><ul style="list-style-type: none"> <li>• false &lt;- Default</li> <li>• true</li> </ul>  |

Before running `omnia.yml`, it is mandatory that red hat subscription be set up on compute nodes running RHEL.

- To set up Red hat subscription, fill in the `rhsm_config.yml` file. Once it's filled in, run the template using Ansible.
- The flow of the playbook will be determined by the value of `redhat_subscription_method` in `rhsm_config.yml`.
  - If `redhat_subscription_method` is set to `portal`, pass the values `username` and `password`. For CLI, run the command:

```
cd utils
ansible-playbook rhsm_subscription.yml -i inventory -e redhat_subscription_
  username="<username>" -e redhat_subscription_password="<password>"
```

- If `redhat_subscription_method` is set to `satellite`, pass the values `organizational identifier` and `activation key`. For CLI, run the command:

```
cd utils
ansible-playbook rhsm_subscription.yml -i inventory -e redhat_subscription_
  activation_key="<activation-key>" -e redhat_subscription_org_id="<org-id>"
```

Where the inventory refers to a file listing all manager and compute nodes per the format provided in [inventory file](#). The inventory file is case-sensitive. Follow the casing provided in the sample file link.

### 3.6.4 Red Hat Unsubscription

To disable subscription on RHEL nodes, the `red_hat_unregister_template` has to be called:

```
cd utils
ansible-playbook rhsm_unregister.yml -i inventory
```

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

### 3.6.5 Set PXE NICs to Static

Use the below playbook to optionally set all PXE NICs on provisioned nodes to 'static'.

**To run the playbook:**

```
cd utils
ansible-playbook configure_pxe_static.yml -i inventory
```

Where inventory refers to a list of IPs separated by newlines:

```
10.5.0.102
10.5.0.103
```

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

### 3.6.6 FreeIPA installation on the NFS node

IPA services are used to provide account management and centralized authentication.

To customize your installation of FreeIPA, enter the following parameters in `input/security_config.yml`.

| Input Parameter      | Definition  | Variable value   |
|----------------------|---|--|
| kerberos_admin       | “admin” user password for the IPA server on RockyOS and RedHat. | The password can be found in the file <code>input/security_config.yml</code> .   |
| ipa_server_hostname  | The hostname of the IPA server                                  | The hostname can be found on the manager node.   |
| domain_name          | Domain name   | The domain name can be found in the file <code>input/security_config.yml</code> .  |
| ipa_server_ipaddress | The IP address of the IPA server                                | The IP address can be found on the IPA server on the manager node using the <code>ip</code> a command. This IP address should be accessible from the NFS node. |

To set up IPA services for the NFS node in the target cluster, run the following command from the `utils/cluster` folder on the control plane:

```
cd utils/cluster
ansible-playbook install_ipa_client.yml -i inventory -e kerberos_admin_password="" -e
ipa_server_hostname="" -e domain_name="" -e ipa_server_ipaddress=""
```

#### Hostname requirements

- The Hostname should not contain the following characters: , (comma), . (period) or \_ (underscore). However, the **domain name** is allowed commas and periods.
- The Hostname cannot start or end with a hyphen (-).
- No upper case characters are allowed in the hostname.
- The hostname cannot start with a number.
- The hostname and the domain name (that is: `hostname00000x.domain.xxx`) cumulatively cannot exceed 64 characters. For example, if the `node_name` provided in `input/provision_config.yml` is ‘node’, and the `domain_name` provided is ‘omnia.test’, Omnia will set the hostname of a target cluster node to ‘node00001.omnia.test’. Omnia appends 6 digits to the hostname to individually name each target node.

**Note:** Use the format specified under [NFS inventory in the Sample Files](#) for inventory.

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

### 3.6.7 Uninstalling the provision tool

Use this script to undo all the changes made by the provision tool. For a list of actions taken by the provision tool, [click here](#).

To run the script:

```
cd utils
ansible-playbook control_plane_cleanup.yml
```

#### Caution:

- When re-provisioning your cluster (that is, re-running the `provision.yml` playbook) after a clean-up, ensure to use a different `admin_nic_subnet` in `input/provision_config.yml` to avoid a conflict with newly assigned servers. Alternatively, disable any OS available in the `Boot Option Enable/Disable` section of your BIOS settings (`BIOS Settings > Boot Settings > UEFI Boot Settings`) on all target nodes.
- On subsequent runs of `provision.yml`, if users are unable to log into the server, refresh the ssh key manually and retry.

```
ssh-keygen -R <node IP>
```

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

### 3.6.8 OS Package Update

To install multiple packages on target nodes in a bulk operation, the `package_update.yml` playbook can be leveraged.

#### Prerequisites

- All target nodes should be running RHEL or Rocky (Versions 8.4, 8.5 or 8.6).
- Download the packages (RPMs) for the target nodes and place them in this folder: `/install/post/otherpkgs/<Provision OS.Version>/x86_64/custom_software/Packages`.

---

**Note:** Do not use ISO files for updates or package installations.

---

- Create a package list by creating the following text file (For packages that are not in RHEL repos, name the file `update.otherpkgs.pkglist`. For OS packages, `xxxx.pkglist`) and place in the default path. For example: `/install/post/otherpkgs/<Provision OS.Version>/x86_64/custom_software/update.otherpkgs.pkglist`:

```
custom_software/<package1>-<version1>
custom_software/<package2>-<version2>
custom_software/<package3>-<version3>
```

To customize the package update, enter the following parameters in `utils/package_update_config.yml`:

| Parameter                                  | Details  |
|--|--|
| <b>os_type</b><br>string Required          | The operating system in use on the target cluster nodes.<br>Choices: <ul style="list-style-type: none"> <li>• <code>rhel</code> &lt;- Default</li> <li>• <code>rocky</code></li> </ul>   |
| <b>os_version</b><br>string Required       | OS version of target nodes in the cluster.<br><b>Default value:</b> 8.6  |
| <b>package_list</b><br>string Required     | <ul style="list-style-type: none"> <li>• Location path for the package list file</li> <li>• For other packagelist, file name should be - (xxx.otherpkgs.pkglist)</li> <li>• For os packagelist, file name should be - (xxx.pkglist)</li> <li>• All packages in this list will be installed/updated on remote nodes</li> </ul> <b>Default value:</b> <code>"/install/post/otherpkgs/rhels8.6.0/x86_64/custom_software/update.otherpkgs.pkglist"</code>  |
| <b>package_type</b><br>string Required     | <ul style="list-style-type: none"> <li>• Indicates whether the packages to be installed are <code>os</code> packages (ie they are available in baseos or appstream) or <code>other</code> (ie they're not part of os repos, appstream or baseos).</li> <li>• If the package is being downloaded to <code>/install/post/otherpkgs/&lt;Provision OS.Version&gt;/x86_64/custom_software/Packages/</code>, use the value <code>other</code>.</li> </ul> Choices: <ul style="list-style-type: none"> <li>• <code>os</code></li> <li>• <code>other</code> &lt;- Default</li> </ul> |
| <b>odelist</b><br>string Required          | Comma separated list of all target nodes in the cluster.<br><b>Default value:</b> <code>all</code>   |
| <b>reboot_required</b><br>boolean Required | Indicates whether the remote nodes listed will be re-booted.<br>Choices: <ul style="list-style-type: none"> <li>• <code>true</code></li> <li>• <code>false</code> &lt;- Default</li> </ul>   |

To run the playbook, run the following commands:

```
cd utils
ansible-playbook package_update.yml
```

**Note:** At the end of the playbook, the package update status is displayed by target node. If the update status of any node is failed, use the command `log (/var/log/xcat/commands.log)` to debug the issue. Alternatively, verify

that the node is reachable post provisioning.

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

### 3.6.9 Clearing ports from Omnia

To undo the configurations made by Omnia to switch ports in the event of a reconfiguration or a clean-up, the `delete_switch_ports.yml` playbook can be utilized.

Enter the required details in `utils/provision/switch_based_deletion_config.yml`:

| Parameter   | Details  |
|---|--|
| <b>switch_based_details</b><br>JSON List Required | <ul style="list-style-type: none"> <li>JSON list of ports to be cleared from the Omnia DBs.</li> <li>Example:               <pre>- { ip: 172.96.28.12, ports: '1-48, ↪49:3,50' }</pre> </li> <li>Example with 2 switches:               <pre>- { ip: 172.96.28.12, ports: '1-48, ↪49:3,50' }  - { ip: 172.96.28.14, ports: '1,2,3,5 ↪' }</pre> </li> </ul> |

To run the playbook, use the below commands:

```
cd utils/provision
ansible-playbook delete_switch_ports.yml
```

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

### 3.6.10 TimescaleDB utility

Telemetry metrics stored in a timescaleDB can be copied locally in a csv format. This file can be used to generate insights into key statistics in your cluster.

To customize the local copy of the timescale DB, fill out the below parameters in `utils/timescaledb_utility/timescaledb_utility_config.yml`.

| Parameter                           | Details  |
|-------------------------------------|--|
| <b>column_name</b> string Optional  | <ul style="list-style-type: none"> <li>Filters timescaleDB data by metric name <b>and</b> value.</li> <li>If this value is not provided, all metrics and their corresponding values will be stored in the file.</li> </ul> |
| <b>column_value</b> string Optional | <ul style="list-style-type: none"> <li>Filters timescaleDB data by metric name <b>and</b> value.</li> <li>If this value is not provided, all metrics and their corresponding values will be stored in the file.</li> </ul> |
| <b>start_time</b> string Optional   | <ul style="list-style-type: none"> <li>Filters timescaleDB data by time of polling.</li> <li>If this value is not provided, all metric values collected will be stored.</li> </ul>   |
| <b>stop_time</b> string Optional    | <ul style="list-style-type: none"> <li>Filters timescaleDB data by time of polling.</li> <li>If this value is not provided, all metric values collected will be stored.</li> </ul>   |
| <b>filename</b> string Required     | <ul style="list-style-type: none"> <li>Target filepath where all timescaleDB will be backed up.</li> </ul>   |

To initiate the backup to local file, run the following ansible playbook:

```
cd utils/timescaledb_utility
ansible-playbook timescaledb_utility.yml
```

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

## 3.7 Telemetry and visualizations

The telemetry feature allows the set up of Omnia telemetry (to poll values from all Omnia provisioned nodes in the cluster) and/or iDRAC telemetry (To poll values from all eligible iDRACs in the cluster). It also installs [Grafana](#) and [Loki](#) as Kubernetes pods.

To initiate telemetry support, fill out the following parameters in `input/telemetry_config.yml`:



Table 1: Parameters

| Parameter  | Details   |
|--|---|
| <b>idrac_telemetry_support</b><br>boolean <sup>1</sup><br>Required           | <ul style="list-style-type: none"> <li>Enables iDRAC telemetry support and visualizations.</li> <li><b>Values:</b></li> </ul> <pre>* false &lt;- Default</pre> <pre>* true</pre> <hr/> <p><b>Note:</b> When <code>idrac_telemetry_support</code> is true, <code>mysqldb_user</code>, <code>mysqldb_password</code> and <code>mysqldb_root_password</code> become mandatory.</p> <hr/>   |
| <b>omnia_telemetry_support</b><br>boolean <sup>Page 130, 1</sup><br>Required | <ul style="list-style-type: none"> <li>Starts or stops Omnia telemetry</li> <li>If <code>omnia_telemetry_support</code> is true, then at least one of <code>collect_regular_metrics</code> or <code>collect_health_check_metrics</code> or <code>collect_gpu_metrics</code> should be true, to collect metrics.</li> <li>If <code>omnia_telemetry_support</code> is false, telemetry acquisition will be stopped.</li> <li><b>Values:</b></li> </ul> <pre>* false &lt;- Default</pre> <pre>* true</pre>   |
| <b>visualization_support</b><br>boolean <sup>Page 130, 1</sup><br>Required   | <ul style="list-style-type: none"> <li>Enables visualizations.</li> <li><b>Values:</b></li> </ul> <pre>* false &lt;- Default</pre> <pre>* true</pre> <hr/> <p><b>Note:</b> When <code>visualization_support</code> is true, <code>grafana_username</code> and <code>grafana_password</code> become mandatory.</p> <hr/>   |
| <b>appliance_k8s_pod_net_cidr</b><br>string<br>Required                      | <ul style="list-style-type: none"> <li>Kubernetes pod network CIDR for appliance k8s network.</li> <li>Make sure this value does not overlap with any of the host networks.</li> <li><b>Default value:</b> "192.168.0.0/16"</li> </ul>  |
| <b>pod_external_ip_start_range</b><br>string<br>Required                     | <ul style="list-style-type: none"> <li>The start of the range that will be used by Load-balancer for assigning IPs to K8s services in admin NIC subnet configured on the control plane.</li> <li>The first and second octets (x,y) are not used/validated by Omnia. These values are internally calculated based on the value of <code>admin_nic_subnet</code> in <code>input/provision_config.yml</code>.</li> <li>If <code>pod_external_ip_start_range:</code> "x.y.240.100" and <code>pod_external_ip_end_range:</code> "x.y.240.105" and</li> </ul> |
| <b>3.7. Telemetry and visualizations</b>                                     | <ul style="list-style-type: none"> <li>If <code>admin_nic_subnet</code> provided in <code>provision_config.yml</code> is 10.5.0.0, <code>pod_external_ip_start_range</code> will be 10.5.240.100 and</li> </ul>   |

Once you have executed `provision.yml` and has also provisioned the cluster, initiate telemetry on the cluster as part of `omnia.yml`, which configures the cluster with scheduler, storage and authentication using the below command.

```
ansible-playbook omnia.yml -i inventory
```

Optionally, you can initiate only telemetry using the below command:

```
ansible-playbook telemetry.yml -i inventory
```

---

**Note:**

- Depending on the type of telemetry initiated, include the following groups in the inventory:
  - `omnia_telemetry`: manager, compute, [optional] login
  - `idrac_telemetry`: idrac
- If you would like a local backup of the timescaleDB used to store telemetry data, [click here](#).

---

After initiation, new iDRACs can be added for `idrac_telemetry` acquisition by running the following commands:

```
ansible-playbook add_idrac_node.yml -i inventory
```

**Modifying telemetry information**

To modify how data is collected from the cluster, modify the variables in `omnia/input/telemetry_config.yml` and re-run the `telemetry.yml` playbook.

- When `omnia_telemetry_support` is set to false, Omnia Telemetry Acquisition service will be stopped on all cluster nodes provided in the passed inventory.
- When `omnia_telemetry_support` is set to true, Omnia Telemetry Acquisition service will be restarted on all cluster nodes provided in the passed inventory.
- To start or stop the collection of regular metrics, health check metrics, or GPU metrics, update the values of `collect_regular_metrics`, `collect_health_check_metrics`, or `collect_gpu_metrics`. For a list of all metrics collected, [click here](#).

---

**Note:**

- Currently, changing the `grafana_username` and `grafana_password` values is not supported via `telemetry.yml`.
- The passed inventory should have an `idrac` group, if `idrac_telemetry_support` is true.
- If `omnia_telemetry_support` is true, then the inventory should have manager and compute groups along with optional login group.
- Rocky 8.7 is not compatible with the Kubernetes installed by `telemetry.yml` due to known issues with cri-o. For more information, [click here](#).
- If a subsequent run of `telemetry.yml` fails, the `telemetry_config.yml` file will be unencrypted.

---

**To access the Grafana UI***Pre requisites*

- `visualisation_support` should be set to true when running `telemetry.yml` or `omnia.yml`.

---

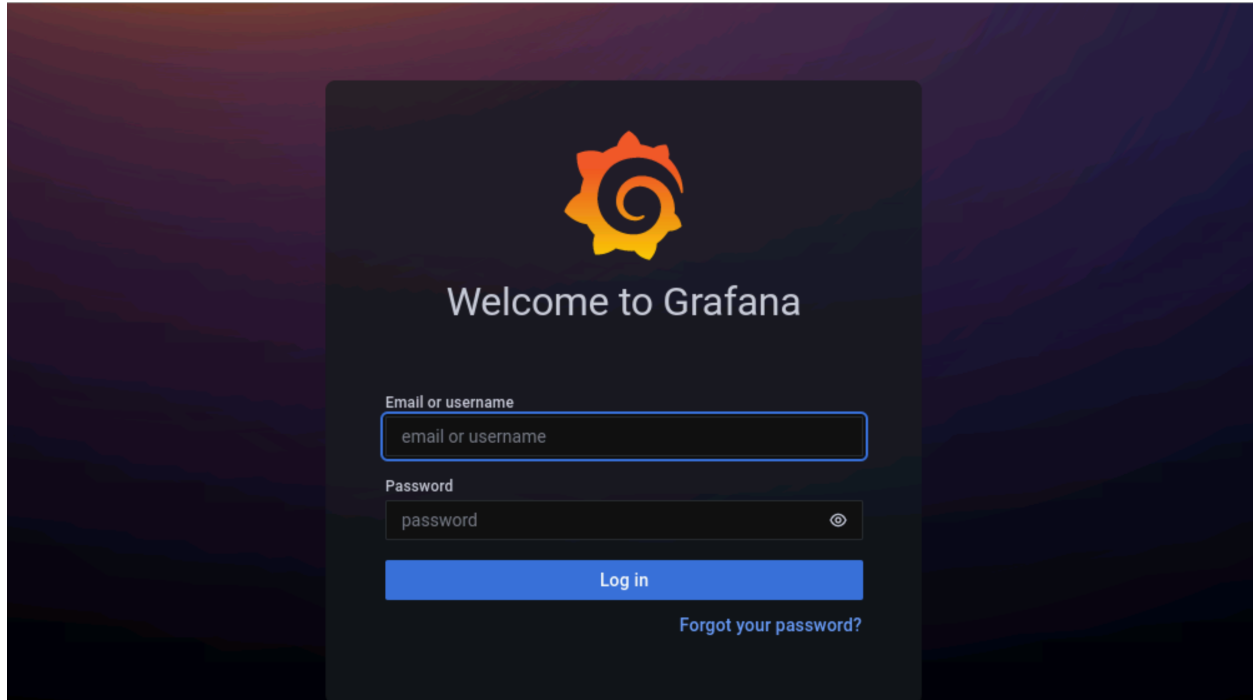
<sup>1</sup> Boolean parameters do not need to be passed with double or single quotes.

- i. Find the IP address of the Grafana service using `kubectl get svc -n grafana`

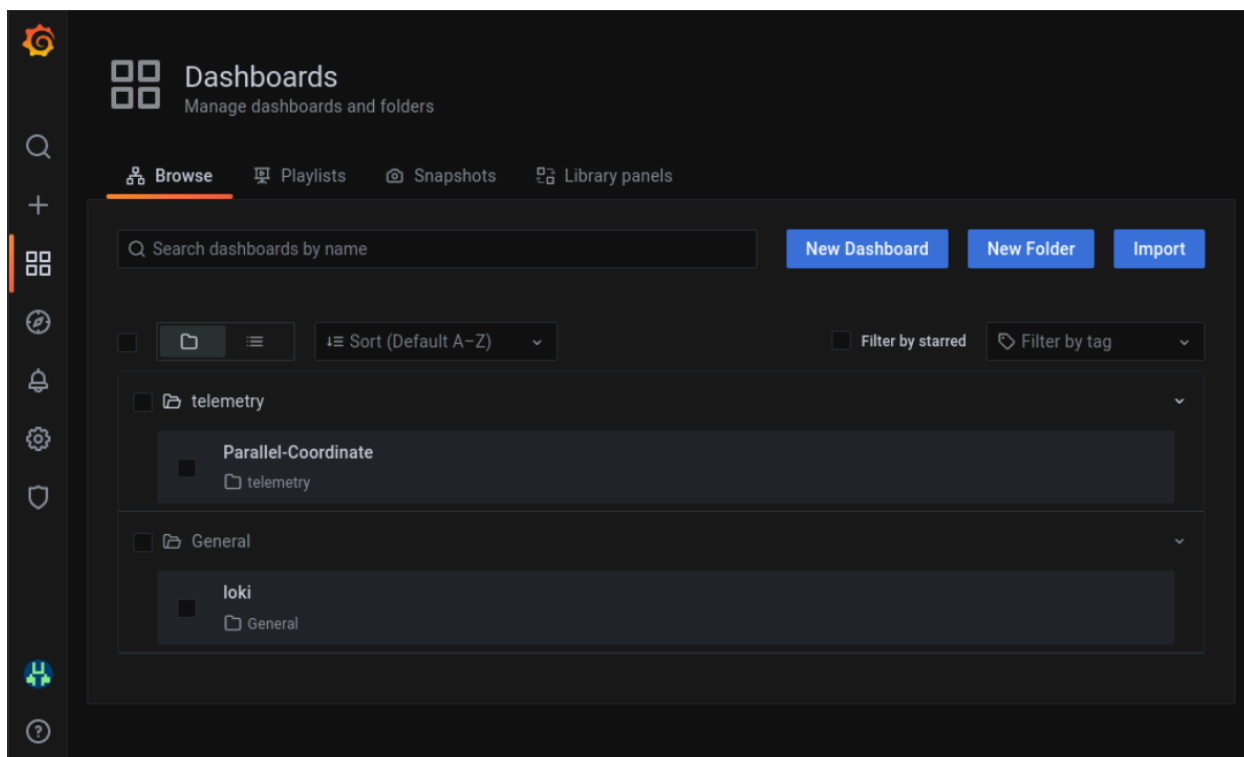
```
[root@localcp ~]# kubectl get svc -n grafana
```

| NAME    | TYPE         | CLUSTER-IP    | EXTERNAL-IP  | PORT(S)        | AGE |
|---------|--------------|---------------|--------------|----------------|-----|
| grafana | LoadBalancer | 10.102.225.70 | 10.2.240.100 | 5000:30590/TCP | 30h |
| loki    | ClusterIP    | 10.100.66.207 | <none>       | 3100/TCP       | 30h |

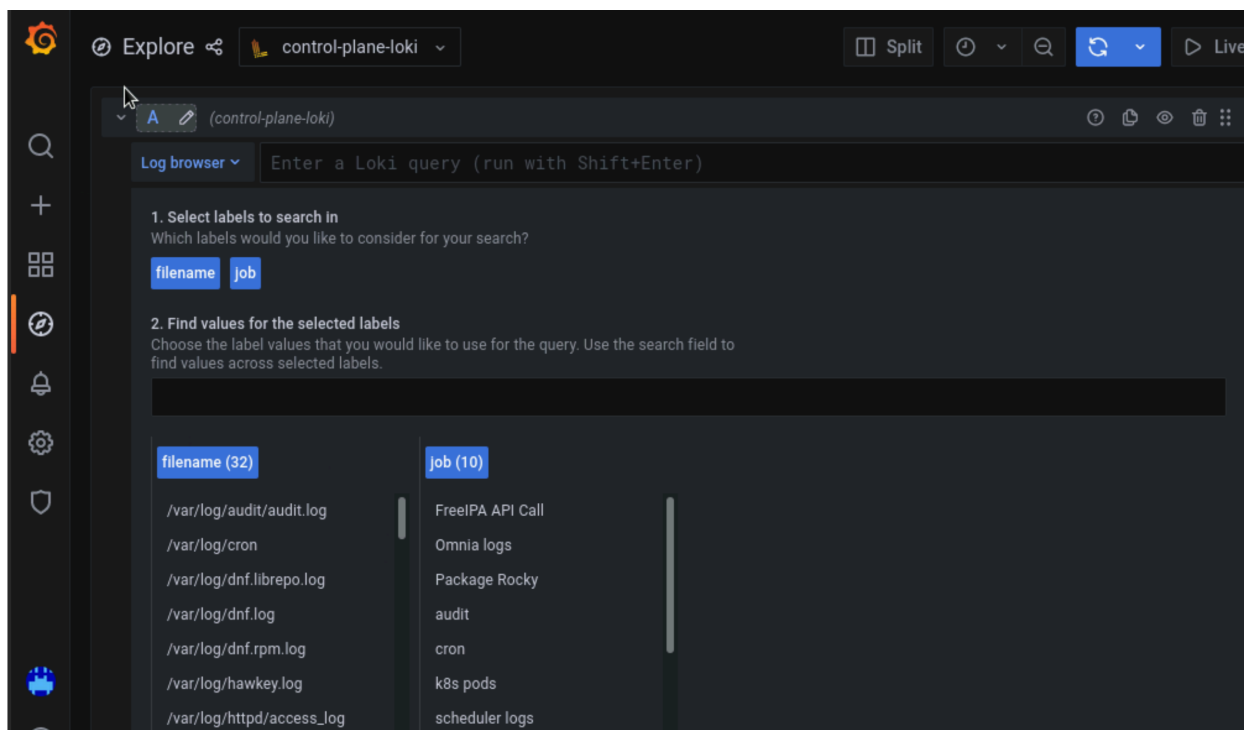
- ii. Login to the Grafana UI by connecting to the cluster IP of grafana service obtained above via port 5000. That is `http://xx.xx.xx.xx:5000/login`



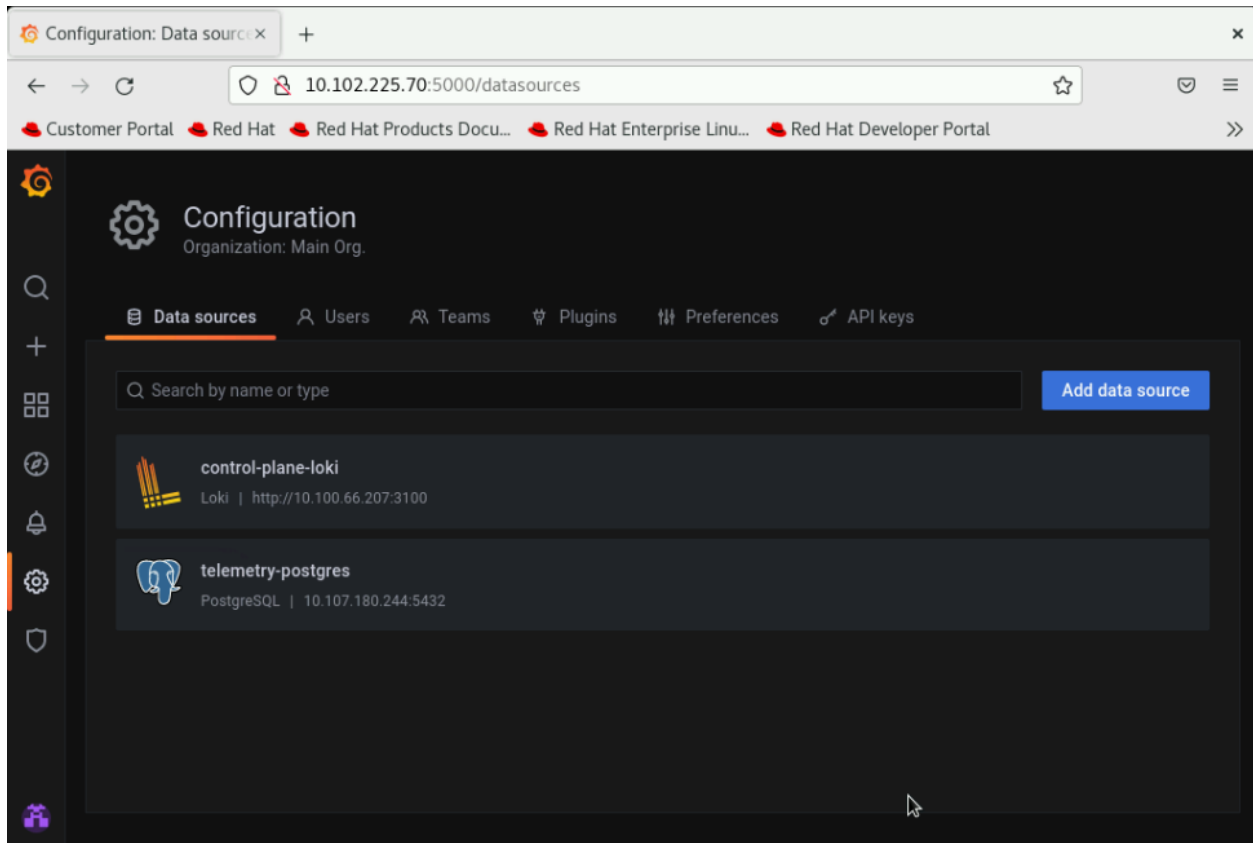
- iii. Enter the `grafana_username` and `grafana_password` as mentioned in `input/telemetry_config.yml`.



Loki log collections can viewed on the explore section of the grafana UI.

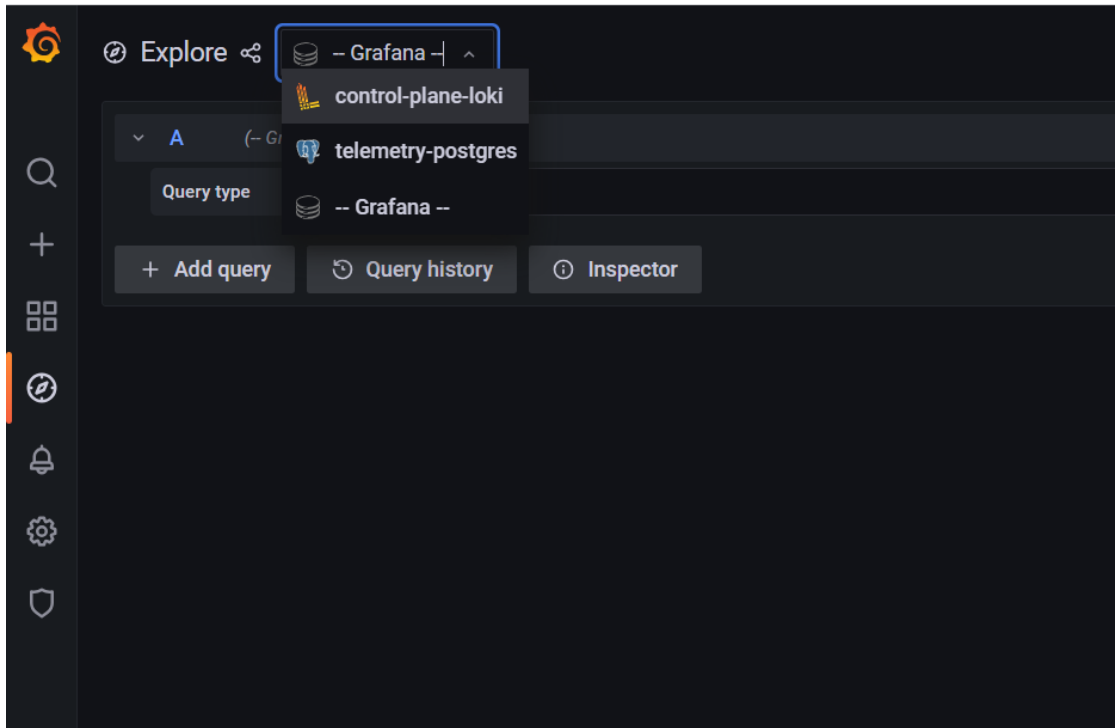


Datasources configured by Omnia can be viewed as seen below.



### To use Loki for log filtering

- i. Login to the Grafana UI by connecting to the cluster IP of grafana service obtained above via port 5000. That is `http://xx.xx.xx.xx:5000/login`
- ii. In the Explore page, select **control-plane-loki**.




iii. The log browser allows you to filter logs by job, node and/or user.

Ex:

```
(job)= "cluster deployment logs") |= "nodename"
(job="compute log messages") |= "nodename" |= "node_username"
```

### To use Grafana to view telemetry data

- i. Login to the Grafana UI by connecting to the cluster IP of grafana service obtained above via port 5000. That is `http://xx.xx.xx.xx:5000/login`
- ii. In the Explore page, select **telemetry-postgres**.



iii. The query builder allows you to create SQL commands that can be used to query the `omnia_telemetry.metrics` table. Filter the data required using the following fields:

- **id**: The name of the metric.
- **context**: The type of metric being collected (Regular Metric, Health Check Metric and GPU metric).

- **label:** A combined field listing the **id** and **context** row values.
- **value:** The value of the metric at the given timestamp.
- **unit:** The unit measure of the metric (eg: Seconds, kb, percent, etc.)
- **system:** The service tag of the cluster node.
- **hostname:** The hostname of the cluster node.
- **time:** The timestamp at which the metric was polled from the cluster node.

If you are more comfortable using SQL queries over the query builder, click on **Edit SQL** to directly provide your query. Optionally, the data returned from a query can be viewed as a graph.

### Visualizations

If `idrac_telemetry_support` and `visualisation_support` is set to true, Parallel Coordinate graphs can be used to view system statistics.

### 3.7.1 Acquiring telemetry data for iDRAC and Omnia

Using tools originally developed by [Texas Tech University's iDVL data visualization lab](#), data polled from iDRAC and Slurm can be processed to generate live graphs. These Graphs can be accessed on the Grafana UI.

Once `provision.yml` is executed and Grafana is set up, use `telemetry.yml` to initiate the Graphs. Data polled via Slurm and iDRAC is streamed into internal databases. This data is processed to create parallel coordinate graphs.

---

**Note:** This feature only works on nodes using iDRACs with a datacenter license running a minimum firmware of 4.0.

---


#### To access the grafana UI:

- i. Find the IP address of the Grafana service using `kubectl get svc -n grafana`

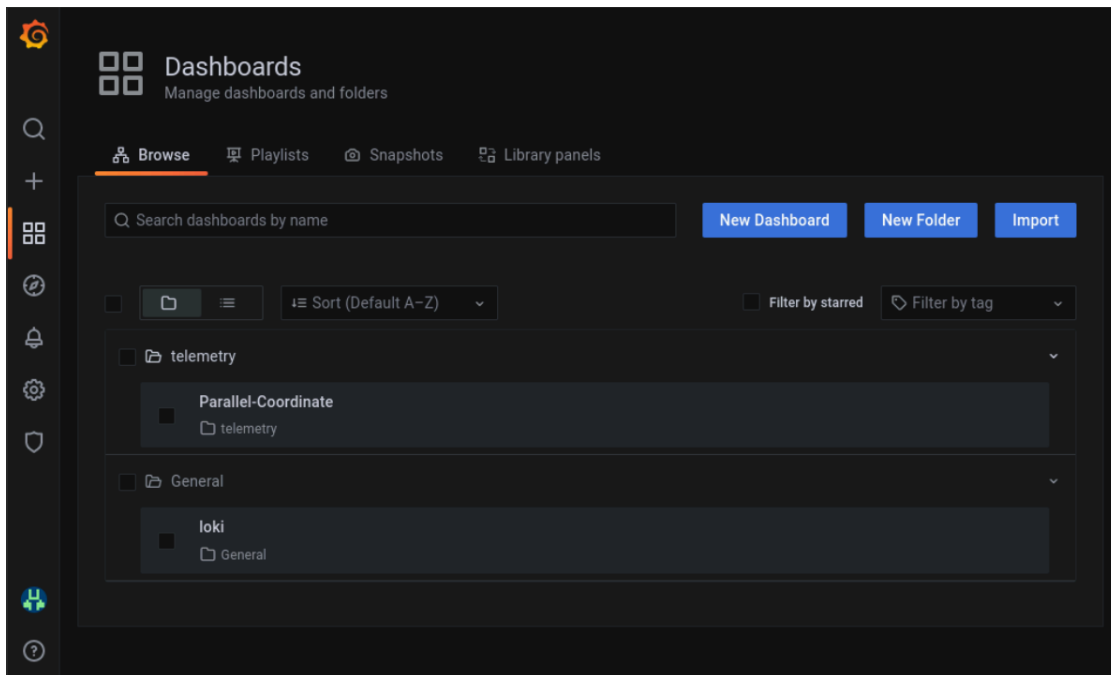
```
[root@localcp ~]# kubectl get svc -n grafana
```

| NAME    | TYPE         | CLUSTER-IP    | EXTERNAL-IP  | PORT(S)        | AGE |
|---------|--------------|---------------|--------------|----------------|-----|
| grafana | LoadBalancer | 10.102.225.70 | 10.2.240.100 | 5000:30590/TCP | 30h |
| loki    | ClusterIP    | 10.100.66.207 | <none>       | 3100/TCP       | 30h |

- ii. Login to the Grafana UI by connecting to the cluster IP of grafana service obtained above via port 5000. That is `http://xx.xx.xx.xx:5000/login`



- iii. Enter the `grafana_username` and `grafana_password` as mentioned in `input/telemetry_config.yml`.



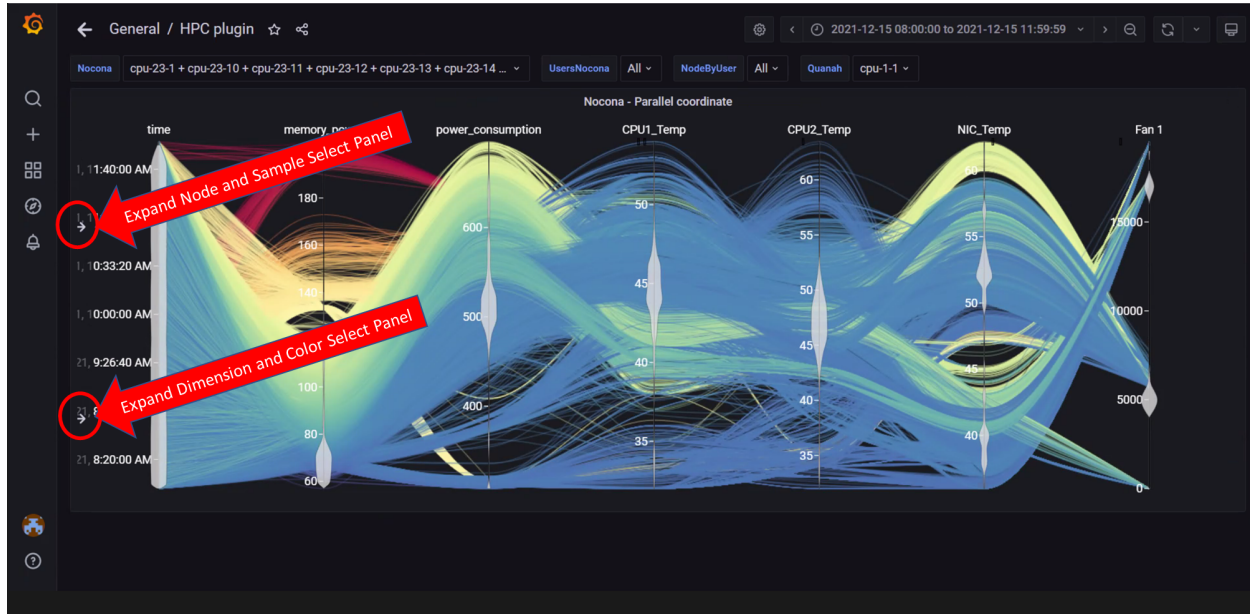
### All your data in a glance:

If `idrac_telemetry_support` and `visualisation_support` is set to true, Parallel Coordinate graphs can be used to view system statistics.

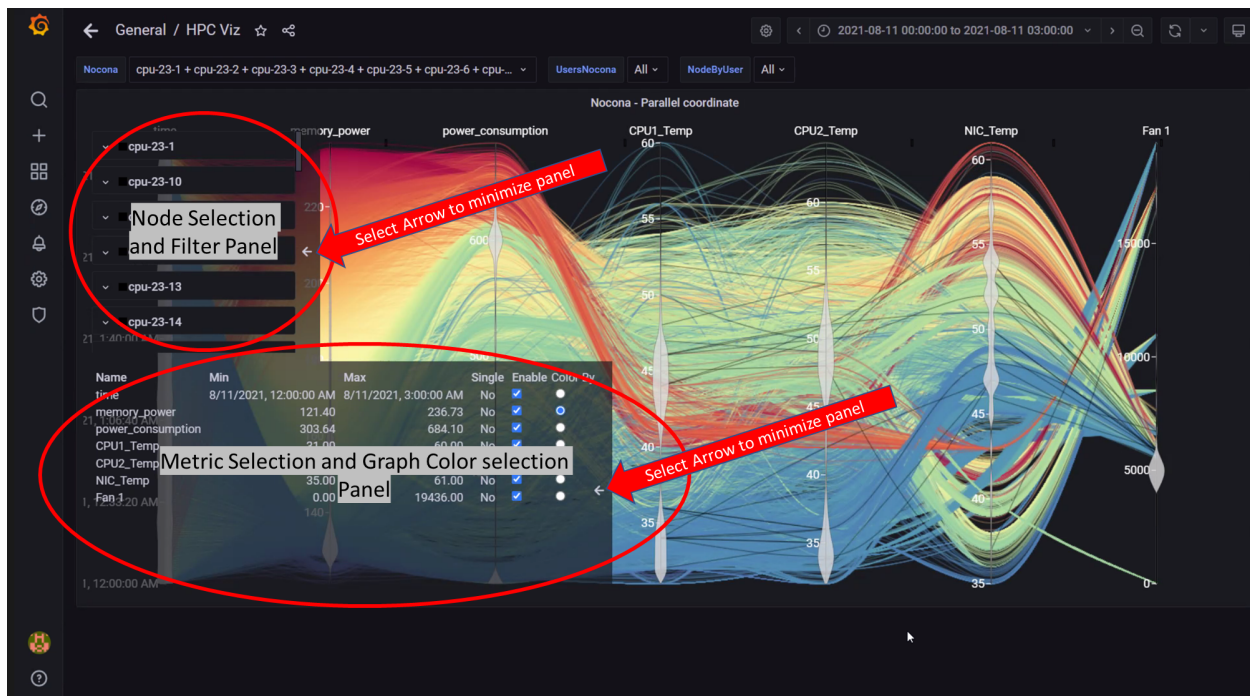


## Parallel coordinates

Parallel coordinates are a great way to visualize multiple metric dimensions simultaneously to see trends and spot outlier activity. Metrics like CPU temp, Fan Speed, Memory Usage etc. can be added or removed as an additional vertical axis. This implementation of parallel coordinate graphing includes a display of metric value distribution in the form of a violin plot along vertical axes and the ability to interact with the graph to perform filtering. Metric range filtering on one or more axes automatically filters the node and sample list in the top left-hand panel to the nodes and samples that fit the filtering criteria.



In the above image, both left-hand panels are collapsed to allow for a better view of the graph. They can be expanded by clicking on the arrows highlighted in the picture. The expanded panels can be used to customize the graph.



In the above image, both left-hand panels are expanded and can be minimized by clicking on the minimize arrows on the right of each panel. These panels can be used to customize the graphs by:

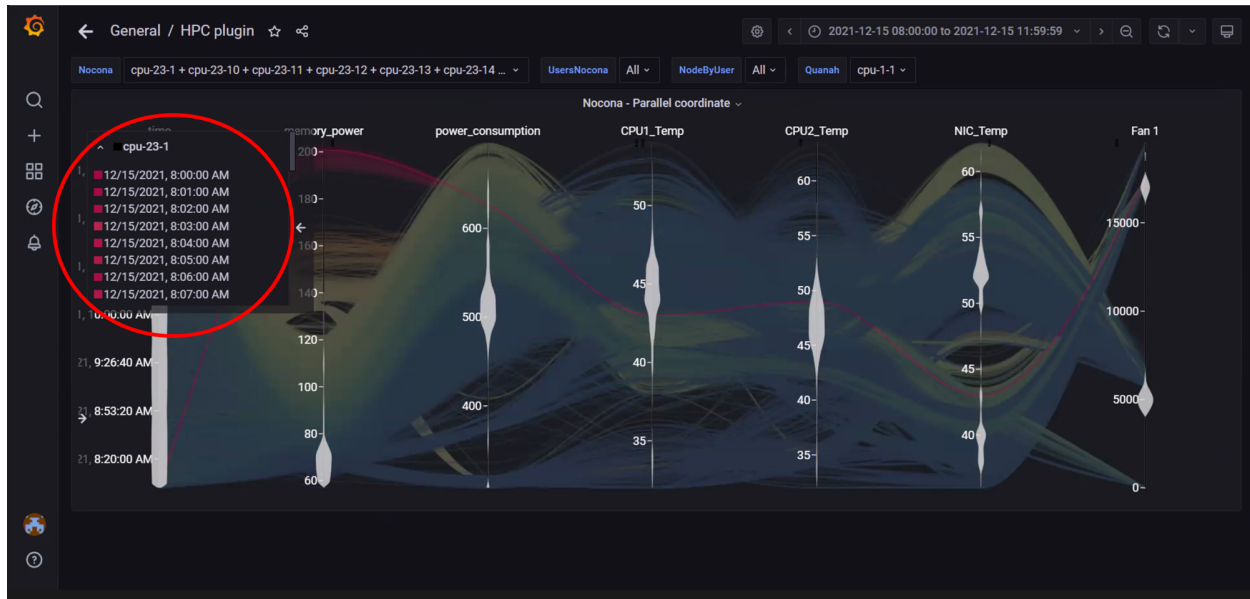
- Filtering by node and node metrics
- Assigning colors to different node metrics



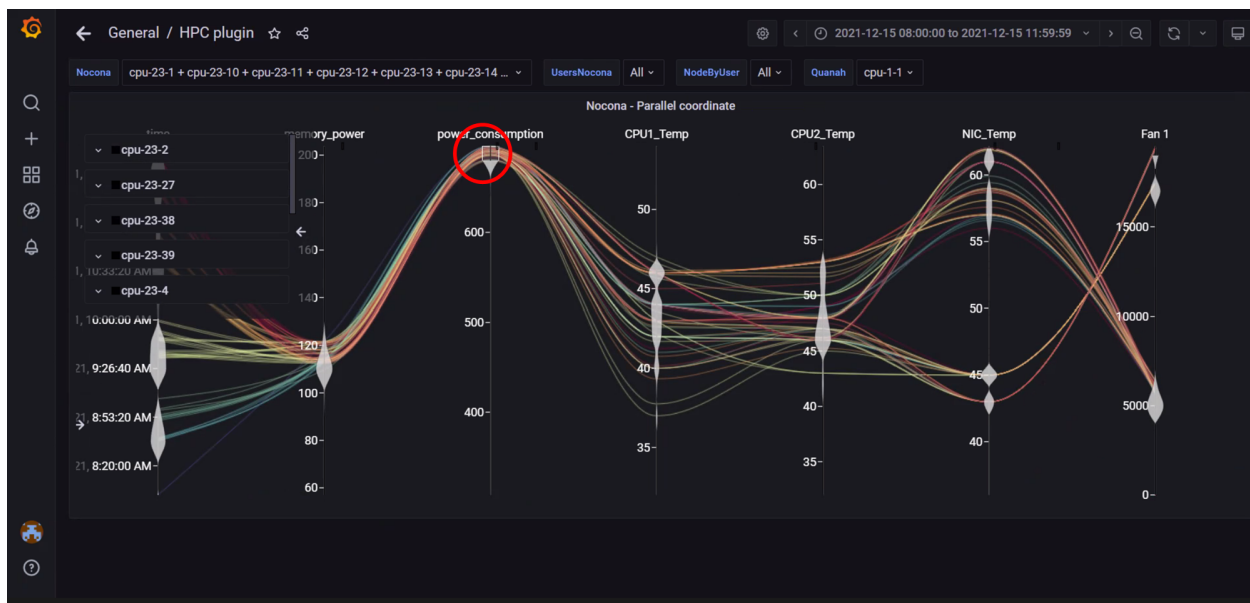
In the above image, the metric **Power Consumption** has been assigned a color to highlight the metric.



In the above image, data has been filtered by **Node** to get insights into different metrics about specific nodes.

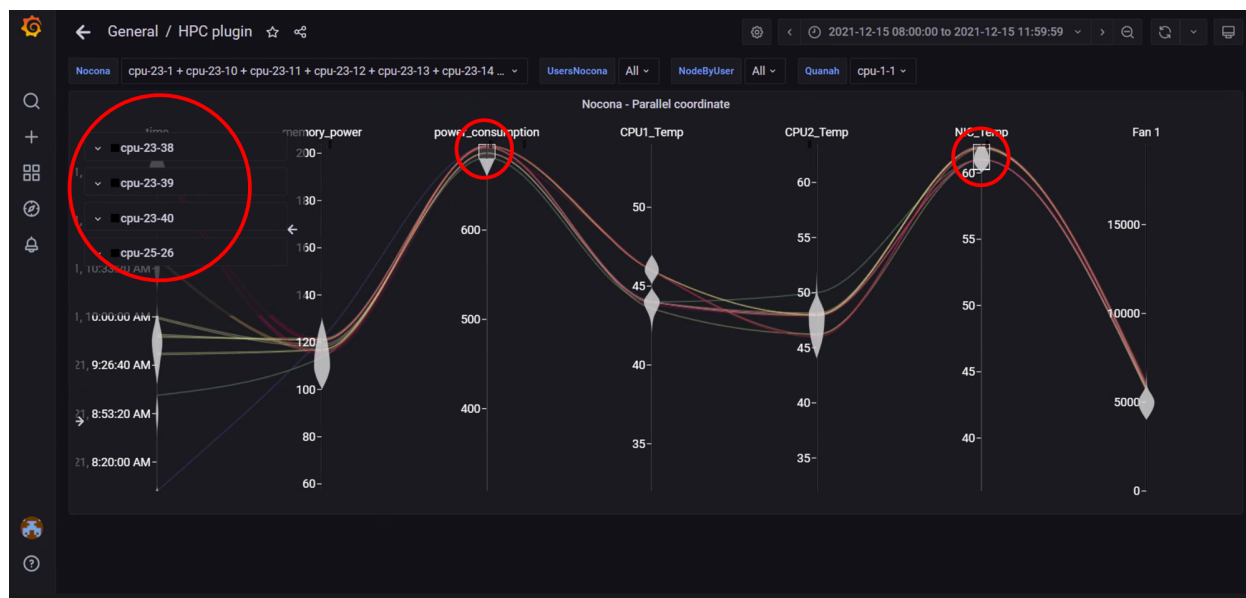


In the above image, data for a single node has been highlighted using the top-left panel.

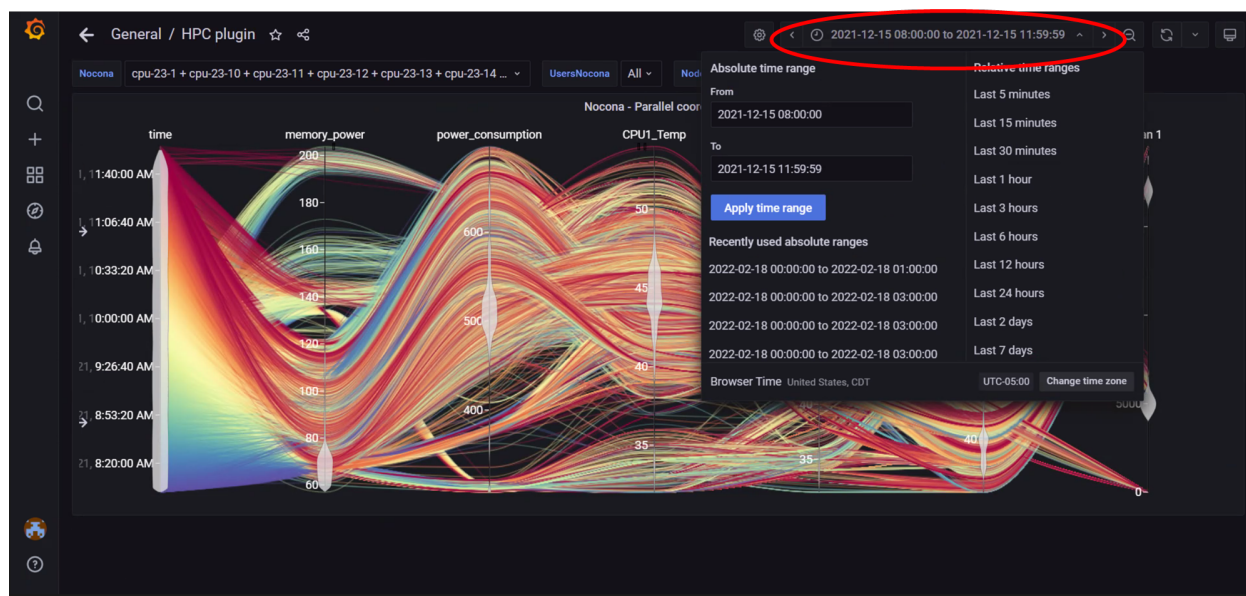


In the above image, metric filters were applied on **Power Consumption** by clicking on the vertical axis and dragging a filter box over the range of values required. The top left panel will display nodes and samples that fit the filter. Filters are removed by clicking on the vertical dimension axis again.





In the above image, metric filters were applied on **Power Consumption** and **NIC temperature**. Using more than one filter will result in fewer nodes and telemetry samples that meet the filtering criteria.



In the above image, the top-right panel was used to filter data by time, this can be done in 2 ways:

- In absolute yyyy-mm-dd hh:mm:ss format
- In relative time periods such as 'last 5 minutes', 'last 7 days' etc.

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

**Note:** The timestamps used for the time metric are based on the timezone set in `input/provision_config.yml`. In the event of a mismatch between the timezone on the browser being used to access Grafana UI and the timezone in `input/provision_config.yml`, the time range being used to filter information on the Grafana UI will have to be adjusted per the timezone in `input/provision_config.yml`.

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

### **3.7.2 Metrics collected**

#### **Regular metrics**

Your cluster in numbers: Regular metrics include information such as CPU, memory, packets errors, drives etc.

Table 2: Regular metrics

| Metric Name             | Unit      | Possible Values   | Possible error causes   |
|-------------------------|-----------|---|---|
| BlockedProcesses        | processes | <ul style="list-style-type: none"> <li>• Metric Value</li> <li>• No Data</li> </ul> | <ul style="list-style-type: none"> <li>• This could happen if the <code>/proc/stat</code> file is inaccessible.</li> </ul>    |
| CPUSystem               | seconds   | <ul style="list-style-type: none"> <li>• Metric Value</li> <li>• No Data</li> </ul> | <ul style="list-style-type: none"> <li>• This could happen when the <code>psutil</code> library encounters errors.</li> </ul> |
| CPUWait                 | seconds   | <ul style="list-style-type: none"> <li>• Metric Value</li> <li>• No Data</li> </ul> | <ul style="list-style-type: none"> <li>• This could happen when the <code>psutil</code> library encounters errors.</li> </ul> |
| ErrorsRecv              |           | <ul style="list-style-type: none"> <li>• Metric Value</li> <li>• No Data</li> </ul> | <ul style="list-style-type: none"> <li>• This could happen when the <code>psutil</code> library encounters errors.</li> </ul> |
| ErrorsSent              |           | <ul style="list-style-type: none"> <li>• Metric Value</li> <li>• No Data</li> </ul> | <ul style="list-style-type: none"> <li>• This could happen when the <code>psutil</code> library encounters errors.</li> </ul> |
| FailedJobs              |           | <ul style="list-style-type: none"> <li>• Metric Value</li> <li>• No Data</li> </ul> | <ul style="list-style-type: none"> <li>• Slurm is not installed.</li> </ul>   |
| HardwareCorruptedMemory | kB        | <ul style="list-style-type: none"> <li>• Metric Value</li> <li>• No Data</li> </ul> | <ul style="list-style-type: none"> <li>• This could happen if the <code>/proc/meminfo</code> file is inaccessible.</li> </ul> |
| MemoryActive            | bytes     | <ul style="list-style-type: none"> <li>• Metric Value</li> <li>• No Data</li> </ul> | <ul style="list-style-type: none"> <li>• This could happen when the <code>psutil</code> library encounters errors.</li> </ul> |
| MemoryAvailable         | bytes     | <ul style="list-style-type: none"> <li>• Metric Value</li> <li>• No Data</li> </ul> | <ul style="list-style-type: none"> <li>• This could happen when the <code>psutil</code> library encounters errors.</li> </ul> |
| MemoryCached            | bytes     | <ul style="list-style-type: none"> <li>• Metric Value</li> <li>• No Data</li> </ul> | <ul style="list-style-type: none"> <li>• This could happen when the <code>psutil</code> library encounters errors.</li> </ul> |
| MemoryFree              | bytes     | <ul style="list-style-type: none"> <li>• Metric Value</li> <li>• No Data</li> </ul> | <ul style="list-style-type: none"> <li>• This could happen when the <code>psutil</code> library encounters errors.</li> </ul> |

## Health metrics

The health of your cluster: Health metrics include key performance indicators.

Table 3: Health metrics

| Metric Name                        | Possible value(s)   | Possible failure causes  |
|------------------------------------|---|--|
| dmesg                              | <ul style="list-style-type: none"> <li>• Unknown</li> <li>• Fail</li> <li>• Pass</li> </ul> | <ul style="list-style-type: none"> <li>• [Unknown] The dmesg command was not found on the cluster node.</li> <li>• [Fail] The dmesg command returned an error log message.</li> </ul>                                      |
| beegfs -beegfsstat                 | <ul style="list-style-type: none"> <li>• Unknown</li> <li>• Fail</li> <li>• Pass</li> </ul> | <ul style="list-style-type: none"> <li>• [Unknown] BeeGFS is not installed or inactive.</li> <li>• [Fail] The BeeGFS client service has failed or the node is not present in reachable lists of BeeGFS clients.</li> </ul> |
| gpu_driver_health:gpu              | <ul style="list-style-type: none"> <li>• Unknown</li> <li>• Fail</li> <li>• Pass</li> </ul> | <ul style="list-style-type: none"> <li>• AMD/NVIDIA accelerators are not present.</li> <li>• GPU drivers are not installed including Rocm and CUDA.</li> </ul>   |
| gpu_health_nvlink:gpu <sup>1</sup> | <ul style="list-style-type: none"> <li>• Unknown</li> <li>• Fail</li> <li>• Pass</li> </ul> | <ul style="list-style-type: none"> <li>• AMD/NVIDIA accelerators are not present.</li> <li>• NVLinks are not present.</li> <li>• GPU drivers are not installed including Rocm and CUDA.</li> </ul>                         |
| gpu_health_pcie:gpu                | <ul style="list-style-type: none"> <li>• Unknown</li> <li>• Fail</li> <li>• Pass</li> </ul> | <ul style="list-style-type: none"> <li>• AMD/NVIDIA accelerators are not present.</li> <li>• GPU drivers are not installed including Rocm and CUDA.</li> </ul>   |
| gpu_health_pmu:gpu                 | <ul style="list-style-type: none"> <li>• Unknown</li> <li>• Fail</li> <li>• Pass</li> </ul> | <ul style="list-style-type: none"> <li>• AMD/NVIDIA accelerators are not present.</li> <li>• GPU drivers are not installed including Rocm and CUDA.</li> </ul>   |
| gpu_health_power:gpu               | <ul style="list-style-type: none"> <li>• Unknown</li> <li>• Fail</li> <li>• Pass</li> </ul> | <ul style="list-style-type: none"> <li>• AMD/NVIDIA accelerators are not present.</li> <li>• GPU drivers are not installed including Rocm and CUDA.</li> </ul>   |
| gpu_health_thermal:gpu             | <ul style="list-style-type: none"> <li>• Unknown</li> <li>• Metric Value</li> </ul>         | <ul style="list-style-type: none"> <li>• AMD/NVIDIA accelerators are not present.</li> <li>• GPU drivers are not installed including Rocm and CUDA.</li> </ul>   |
| Kubernetespodsstatus               | <ul style="list-style-type: none"> <li>• Unknown</li> <li>• Fail</li> <li>• Pass</li> </ul> | <ul style="list-style-type: none"> <li>• Kubernetes is not installed.</li> </ul>   |



## GPU metrics

The GPUs of your cluster: GPU metrics include information about GPUs in the cluster

Table 4: GPU metrics

| Metric Name             | Unit    | Possible value(s)   | Potential error cause(s)   |
|-------------------------|---------|---|--|
| gpu_temperature:gpu     | C       | <ul style="list-style-type: none"> <li>• Metric value</li> <li>• No data</li> </ul> | <ul style="list-style-type: none"> <li>• AMD/NVIDIA accelerators are not present.</li> <li>• GPU drivers are not installed including Rocm and CUDA.</li> </ul> |
| gpu_utilization         | percent | <ul style="list-style-type: none"> <li>• Metric value</li> <li>• No data</li> </ul> | <ul style="list-style-type: none"> <li>• AMD/NVIDIA accelerators are not present.</li> <li>• GPU drivers are not installed including Rocm and CUDA.</li> </ul> |
| gpu_utilization:average | percent | <ul style="list-style-type: none"> <li>• Metric value</li> <li>• No data</li> </ul> | <ul style="list-style-type: none"> <li>• AMD/NVIDIA accelerators are not present.</li> <li>• GPU drivers are not installed including Rocm and CUDA.</li> </ul> |

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@del.com](mailto:omnia.readme@del.com).

### 3.7.3 Additional metric information

Table 5: Telemetry metrics

| Metric Name      | Command  | Comments  | Aggregation Level |
|------------------|--|---|-------------------|
| BlockedProcesses | grep procs_blocked<br>/proc/stat                                   |   | Node Level        |
| CPUSystem        | psutil.cpu_times().<br>system                                      |   | Node Level        |
| CPUWait          | psutil.cpu_times().<br>iowait                                      |   | Node Level        |
| ErrorsRecv       | psutil.<br>net_io_counters(perni-<br>get(interface_name).<br>errin | Error packets received for<br>individual network inter-<br>faces will be populated. | Node Level        |

continues on next page

<sup>1</sup> While NVLink metrics are collected as part of our telemetry suite, NVLinks have not been tested for compatibility with Omnia.

Table 5 – continued from previous page

| Metric Name                  | Command   | Comments  | Aggregation Level |
|------------------------------|---|---|-------------------|
| ErrorsSent                   | psutil.<br>net_io_counters(perni-<br>get(interface_name).<br>errout | Error packets sent for indi-<br>vidual network interfaces<br>will be populated. | Node Level        |
| FailedJobs                   | sacct -P<br>--delimiter=\t  |   | Cluster Level     |
| HardwareCorruptedMem-<br>ory | grep<br>HardwareCorrupted<br>/proc/meminfo                          |   | Node Level        |
| MemoryActive                 | psutil.<br>virtual_memory().<br>active                              |   | Node Level        |
| MemoryAvailable              | psutil.<br>virtual_memory().<br>available                           |   | Node Level        |
| MemoryCached                 | psutil.<br>virtual_memory().<br>cached                              |   | Node Level        |
| MemoryFree                   | psutil.<br>virtual_memory().<br>free                                |   | Node Level        |
| MemoryInactive               | psutil.<br>virtual_memory().<br>inactive                            |   | Node Level        |
| MemoryPercent                | psutil.<br>virtual_memory().<br>percent                             |   | Node Level        |
| MemoryShared                 | psutil.<br>virtual_memory().<br>shared                              |   | Node Level        |
| MemoryTotal                  | psutil.<br>virtual_memory().<br>total                               |   | Node Level        |
| MemoryUsed                   | psutil.<br>virtual_memory().<br>used                                |   | Node Level        |

continues on next page

Table 5 – continued from previous page

| Metric Name  | Command   | Comments   | Aggregation Level |
|--------------|---|--|-------------------|
| NodesDown    | <code>sinfo --format=%N\ t%P\t%a\t%C\t%t\t%D\ t%m</code>  | Node is considered down if node state is any of the following: down, drained, draining, fail, failing, future, inval, maint, powered_down, powering_down, unknown, unk.<br><br><b>Note:</b> Node state with * in suffix will be considered as down. Example, idle* will be considered as down. | Cluster Level     |
| NodesTotal   | <code>sinfo --format=%N\ t%P\t%a\t%C\t%t\t%D\ t%m</code>  |  | Cluster Level     |
| NodesUp      | <code>sinfo --format=%N\ t%P\t%a\t%C\t%t\t%D\ t%m</code>  | Node is considered up if node state is any of the following: idle, mixed, completing.<br><br><b>Note:</b> Node state with * in suffix will be considered as down node. Example, idle* will be considered as down node.   | Cluster Level     |
| QueuedJobs   | <code>squeue --format=%i\ t%P\t%j\t%u\t%T\t%S\ t%N</code> |  | Cluster Level     |
| RunningJobs  | <code>squeue --format=%i\ t%P\t%j\t%u\t%T\t%S\ t%N</code> |  | Cluster Level     |
| SMARTHDATemp | <code>smartctl -a &lt;device name&gt;</code>              |  | Node Level        |

continues on next page

Table 5 – continued from previous page

| Metric Name                        | Command   | Comments  | Aggregation Level                                       |
|------------------------------------|---|---|---|
| UniqueUserLogin                    | <code>who cut -f 1 -d " " sort -u wc -l</code>  | <ul style="list-style-type: none"> <li>Locally created users via <code>useradd</code> command are also counted in <code>UniqueUserLogin</code> count.</li> <li>Remote logged in LDAP users are not counted in <code>UniqueUserLogin</code> on login nodes.</li> <li>Remote logged in FreeIPA users are counted in <code>UniqueUserLogin</code> on login nodes.</li> </ul> | Login Node/ Manager Node (If Login Node is not present) |
| dmesg                              | <code>dmesg --level=err</code>  |   | Node Level  |
| Beegfs-beegfsstat                  | <code>systemctl is-active beegfs-client beegfs-ctl --nodetype=client --listnodes</code>   |   | Node Level  |
| gpu_driver_health:gpu              | <ul style="list-style-type: none"> <li>For NVIDIA GPU: <code>nvidia-smi --query-gpu=driver_version --format=csv, noheader</code></li> <li>For AMD GPU: <code>rocm-smi --showdriverversion --csv</code></li> </ul> |   | Node Level  |
| gpu_health_nvlink:gpu <sup>1</sup> | NVIDIA: <code>nvidia-smi nvlink --status</code>   |   | Node Level  |
| gpu_health_pcie:gpu                | <ul style="list-style-type: none"> <li>For NVIDIA GPU: <code>nvidia-smi --query-gpu=pci_bus_id --format=csv, noheader</code></li> <li>For AMD GPU: <code>rocm-smi --showbus --csv</code></li> </ul>               |   | Node Level  |

continues on next page

Table 5 – continued from previous page

| Metric Name                | Command   | Comments  | Aggregation Level |
|----------------------------|---|---|-------------------|
| gpu_health_pmu:gpu         | For NVIDIA GPU:<br>nvidia-smi --query-gpu=power.management --format=csv,nounits   | PMU - Power management unit   | Node Level        |
| gpu_health_power:gpu       | For NVIDIA GPU:<br>nvidia-smi --query-gpu=pci.bus_id --format=csv,nounits   | Power consumption   | Node Level        |
| gpu_health_thermal:gpu     | For AMD GPU:<br>rocm-smi --showbus --csv  | GPU temperature health  | Node Level        |
| Kubernetespodsstatus       | sudo kubectl get pods -A -o json  | Value is pass when all pods and containers are in running state, otherwise Fail.                        | Cluster Level     |
| Kuberneteschildnode        | sudo kubectl get nodes -o json  | Value is pass when all child nodes are in Ready or Ready,SchedulingDisabled state, otherwise Fail.      | Cluster Level     |
| kubernetesnodesstatus      | sudo kubectl get nodes -o json  | Value is pass when all nodes are in Ready or Ready,SchedulingDisabled state, otherwise Fail.            | Cluster Level     |
| kubernetescomponentsstatus | sudo kubectl get --raw=/livez?verbose   | Value is Pass when health check is passed in kubectl get --raw=/livez?verbose command , otherwise fail. | Cluster Level     |
| Smart                      | smartctl -a <device name>   |   | Node Level        |
| gpu_temperature:gpu        | <ul style="list-style-type: none"> <li>For NVIDIA GPU: nvidia-smi --query-gpu=temperature.gpu --format=csv,nounits</li> <li>For AMD GPU: rocm-smi -t --csv</li> </ul> |   | Node Level        |

continues on next page

Table 5 – continued from previous page

| Metric Name             | Command   | Comments  | Aggregation Level |
|-------------------------|---|---|-------------------|
| gpu_utilization:        | <ul style="list-style-type: none"> <li>For NVIDIA GPU: <code>nvidia-smi nvidia-smi --query-gpu=utilization --format=csv, nounits</code></li> <li>For AMD GPU: <code>rocm-smi -u --csv</code></li> </ul>   |   | Node Level        |
| gpu_utilization:average | <ul style="list-style-type: none"> <li>*For NVIDIA GPU: <code>nvidia-smi nvidia-smi --query-gpu=utilization --format=csv, nounits</code></li> <li>*For AMD GPU: <code>rocm-smi -u --csv</code></li> </ul> | Value is average of utilization value of all GPUs | Node Level        |

**Note:** psutil (python system and process utilities) is a cross-platform library for retrieving information on running processes and system utilization (CPU, memory, network).

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

### 3.7.4 Timescale DB

#### Accessing the timescale DB

1. Check the IP of the control plane (ifconfig):

```
3: eno8403: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP group
↪ default qlen 1000 link/ether b4:45:06:eb:da:4e brd ff:ff:ff:ff:ff:ff
inet 198.168.0.11/24 brd 198.168.0.255 scope global dynamic noprefixroute eno8403
↪ validlft 30884289sec preferred_lft 30884289sec inet6 fe80::b645:6ff:feeb:da4e/64
↪ scope link noprefixroute validlft forever preferredlft forever
```

2. Check the external port on which timescaleDB is running (`kubectl get svc -A`):

```
[root@orchidcp xcat]# kubectl get svc -A
NAMESPACE   NAME                                     TYPE        CLUSTER-IP   EXTERNAL-IP   PORT(S)          AGE
calico-apiserver   calico-api                             ClusterIP    10.101.90.148 <none>        443/TCP          13h
calico-kube-system calico-kube-controllers-metrics        ClusterIP    None          <none>          9094/TCP         13h
calico-system      calico-typha                           ClusterIP    10.98.124.122 <none>        5473/TCP         13h
default           kubernetes                             ClusterIP    10.96.0.1     <none>        443/TCP          13h
grafana            grafana                                LoadBalancer 10.104.75.196 10.5.240.100 5000:30110/TCP   13h
loki               loki                                   ClusterIP    10.98.233.254 <none>        3100/TCP         13h
kube-system        kube-dns                               ClusterIP    10.96.0.10    <none>        53/UDP,53/TCP,9153/TCP 13h
kubernetes-dashboard dashboard-metrics-scraper               ClusterIP    10.106.196.186 <none>        8000/TCP         13h
kubernetes-dashboard dashboard                               ClusterIP    10.100.31.32   <none>        443/TCP          13h
metallb-system     webhook-service                         ClusterIP    10.99.39.243   <none>        443/TCP          13h
telemetry-and-visualizations mysqldb                                 ClusterIP    10.103.156.80 <none>        3306/TCP,33060/TCP 13h
telemetry-and-visualizations timescaledb                             LoadBalancer 10.104.33.121 10.5.240.101 5432:52160/TCP   13h
```

<sup>1</sup> While NVLink metrics are collected as part of our telemetry suite, NVLinks have not been tested for compatibility with Omnia.

---

Eg:

| telemetry.metrics=# select * from omnia telemetry.metrics; |                     |  |              |           |         |                      |                        |  |  |
|--|---------------------|--|--------------|-----------|---------|----------------------|------------------------|--|--|
| id   | context             | label                                  | value        | unit      | system  | hostname             | time                   |  |  |
| BlockedProcesses   | Regular Metric      | BlockedProcesses Regular Metric        | 0            | processes | 4FCNDX2 | node00003.omnia.test | 2023-10-09 11:12:55+00 |  |  |
| CPUSystem  | Regular Metric      | CPUSystem Regular Metric               | 14051.36     | seconds   | 4FCNDX2 | node00003.omnia.test | 2023-10-09 11:12:55+00 |  |  |
| CPUWait  | Regular Metric      | CPUWait Regular Metric                 | 67.17        | seconds   | 4FCNDX2 | node00003.omnia.test | 2023-10-09 11:12:55+00 |  |  |
| ErrorsRecv:eno16   | Regular Metric      | ErrorsRecv:eno16 Regular Metric        | 0            |           | 4FCNDX2 | node00003.omnia.test | 2023-10-09 11:12:55+00 |  |  |
| ErrorsSent:eno16   | Regular Metric      | ErrorsSent:eno16 Regular Metric        | 0            |           | 4FCNDX2 | node00003.omnia.test | 2023-10-09 11:12:55+00 |  |  |
| HardwareCorruptedMemory                                    | Regular Metric      | HardwareCorruptedMemory Regular Metric | 0            | KB        | 4FCNDX2 | node00003.omnia.test | 2023-10-09 11:12:55+00 |  |  |
| MemoryFree   | Regular Metric      | MemoryFree Regular Metric              | 18992641868  | bytes     | 4FCNDX2 | node00003.omnia.test | 2023-10-09 11:12:55+00 |  |  |
| MemoryTotal  | Regular Metric      | MemoryTotal Regular Metric             | 208952127488 | bytes     | 4FCNDX2 | node00003.omnia.test | 2023-10-09 11:12:55+00 |  |  |
| MemoryAvailable  | Regular Metric      | MemoryAvailable Regular Metric         | 194328633344 | bytes     | 4FCNDX2 | node00003.omnia.test | 2023-10-09 11:12:55+00 |  |  |
| MemoryPercent  | Regular Metric      | MemoryPercent Regular Metric           | 3.3          | percent   | 4FCNDX2 | node00003.omnia.test | 2023-10-09 11:12:55+00 |  |  |
| MemoryUsed   | Regular Metric      | MemoryUsed Regular Metric              | 2352537376   | bytes     | 4FCNDX2 | node00003.omnia.test | 2023-10-09 11:12:55+00 |  |  |
| MemoryActive   | Regular Metric      | MemoryActive Regular Metric            | 2805833728   | bytes     | 4FCNDX2 | node00003.omnia.test | 2023-10-09 11:12:55+00 |  |  |
| MemoryInactive   | Regular Metric      | MemoryInactive Regular Metric          | 7034634240   | bytes     | 4FCNDX2 | node00003.omnia.test | 2023-10-09 11:12:55+00 |  |  |
| MemoryCached   | Regular Metric      | MemoryCached Regular Metric            | 9569834240   | bytes     | 4FCNDX2 | node00003.omnia.test | 2023-10-09 11:12:55+00 |  |  |
| MemoryShared   | Regular Metric      | MemoryShared Regular Metric            | 2927841280   | bytes     | 4FCNDX2 | node00003.omnia.test | 2023-10-09 11:12:55+00 |  |  |
| SMARTAtTemp:/dev/sda                                       | Regular Metric      | SMARTAtTemp:/dev/sda Regular Metric    | 3            | C         | 4FCNDX2 | node00003.omnia.test | 2023-10-09 11:12:55+00 |  |  |
| SMARTHdATemp:/dev/bus/0                                    | Regular Metric      | SMARTHdATemp:/dev/bus/0 Regular Metric | No data      |           | 4FCNDX2 | node00003.omnia.test | 2023-10-09 11:12:55+00 |  |  |
| dmesg  | Health Check Metric | dmesg Health Check Metric              | Fail         |           | 4FCNDX2 | node00003.omnia.test | 2023-10-09 11:12:55+00 |  |  |
| Bogfs -bogfsfstest   | Health Check Metric | Bogfs -bogfsfstest Health Check Metric | Unknown      |           | 4FCNDX2 | node00003.omnia.test | 2023-10-09 11:12:55+00 |  |  |
| Smart:/dev/sda   | Health Check Metric | Smart:/dev/sda Health Check Metric     | Pass         |           | 4FCNDX2 | node00003.omnia.test | 2023-10-09 11:12:55+00 |  |  |
| Smart:/dev/bus/0   | Health Check Metric | Smart:/dev/bus/0 Health Check Metric   | Unknown      |           | 4FCNDX2 | node00003.omnia.test | 2023-10-09 11:12:55+00 |  |  |
| gpu_health_driver  | Health Check Metric | gpu_health_driver Health Check Metric  | Unknown      |           | 4FCNDX2 | node00003.omnia.test | 2023-10-09 11:12:55+00 |  |  |
| gpu_health_nvlink  | Health Check Metric | gpu_health_nvlink Health Check Metric  | Unknown      |           | 4FCNDX2 | node00003.omnia.test | 2023-10-09 11:12:55+00 |  |  |
| gpu_health_pcie  | Health Check Metric | gpu_health_pcie Health Check Metric    | Unknown      |           | 4FCNDX2 | node00003.omnia.test | 2023-10-09 11:12:55+00 |  |  |
| gpu_health_pmu   | Health Check Metric | gpu_health_pmu Health Check Metric     | Unknown      |           | 4FCNDX2 | node00003.omnia.test | 2023-10-09 11:12:55+00 |  |  |
| gpu_health_power   | Health Check Metric | gpu_health_power Health Check Metric   | Unknown      |           | 4FCNDX2 | node00003.omnia.test | 2023-10-09 11:12:55+00 |  |  |
| gpu_health_thermal   | Health Check Metric | gpu_health_thermal Health Check Metric | Unknown      |           | 4FCNDX2 | node00003.omnia.test | 2023-10-09 11:12:55+00 |  |  |
| gpu_temperature  | GPU Metric          | gpu_temperature GPU Metric             | No data      |           | 4FCNDX2 | node00003.omnia.test | 2023-10-09 11:12:55+00 |  |  |
| gpu_utilization  | GPU Metric          | gpu_utilization GPU Metric             | No data      |           | 4FCNDX2 | node00003.omnia.test | 2023-10-09 11:12:55+00 |  |  |
| gpu_utilization:average                                    | GPU Metric          | gpu_utilization:average GPU Metric     | No data      |           | 4FCNDX2 | node00003.omnia.test | 2023-10-09 11:12:55+00 |  |  |

[illegible]

The `omnia_telemetry.metrics` has a data retention policy that ensures data is stored for 2 months only. A cleanup job is run everyday to purge metrics older than 60 days.

### 3.7.5 Performance stats of telemetry acquisition service

Omnia telemetry testing is conducted on a cluster under 3 different conditions:

1. Omnia Telemetry service running on all nodes provided in the inventory file before running MPI jobs.
2. Omnia Telemetry service running on all nodes provided in the inventory file while running MPI jobs.
3. Omnia Telemetry service running on all nodes provided in the inventory file post MPI job execution.

For this performance testing, the below cluster and parameters are considered:

|                                  |   |
|----------------------------------|---|
| <b>Cluster</b>                   | 1 manager node, 1 login node, 3 compute nodes |
| <b>Omnia_collection_interval</b> | 60 seconds                                    |
| <b>Fuzzy_offset</b>              | 60  |
| <b>Metric_collection_timeout</b> | 5 seconds                                     |
| <b>Scheduler_type</b>            | Slurm and Kubernetes                          |

With the above parameters, CPU and memory utilization on nodes are as below:

|                      | Pre MPI jobs                        | While running MPI jobs              | Post MPI jobs                       |
|----------------------|-------------------------------------|-------------------------------------|-------------------------------------|
| <b>Compute node</b>  | Avg Cpu Utilization: 24.05%         | Avg Cpu Utilization: 26.34%         | Avg Cpu Utilization: 31.63%         |
|                      | Average memory utilization: 7.7674% | Average memory utilization: 2.903%  | Average memory utilization: 2.910%  |
| <b>Login node</b>    | Max number of processes: 5          | Max number of processes: 6          | Max number of processes: 5          |
|                      | Avg Cpu Utilization: 24.42%         | Avg Cpu Utilization: 0.37%          | Avg Cpu Utilization: 0.35%          |
|                      | Average memory utilization: 99.98%  | Average memory utilization: 2.6%    | Average memory utilization: 2.65    |
| <b>Manager Node</b>  | Max number of processes: 6          | Max number of processes: 5          | Max number of processes: 5          |
|                      | Avg Cpu Utilization: 26.14%         | Avg Cpu Utilization: 27.03%         | Avg Cpu Utilization: 29.22%         |
|                      | Average memory utilization: 14.307% | Average memory utilization: 15.572% | Average memory utilization: 15.500% |
| <b>Control Plane</b> | Max number of processes: 7          | Max number of processes: 7          | Max number of processes: 8          |
|                      | Avg Cpu Utilization: 3.17%          | Avg Cpu Utilization: 3.27%          | Avg Cpu Utilization: 3.19%          |
|                      | Average memory utilization: 18.68%  | Average memory utilization: 24.85%  | Average memory utilization: 24.93%  |
| <b>Overall</b>       | Avg Cpu Utilization: 24.87%         | Avg Cpu Utilization: 84.58%         | Avg Cpu Utilization: 20.40%         |
|                      | Average memory utilization: 7.69%   | Average memory utilization: 7.02%   | Average memory utilization: 7.003%  |
|                      | Max number of processes: 7          | Max number of processes: 7          | Max number of processes: 8          |

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).



## LOGGING

### 4.1 Log management

Use `/etc/logrotate.conf` to customize how often logs are rotated. By default, the `logrotate.conf` file:

```
cat /etc/logrotate.conf
# see "man logrotate" for details
# rotate log files weekly
weekly
# keep 4 weeks worth of backlogs
rotate 4
# create new (empty) log files after rotating old ones
create
# use date as a suffix of the rotated file
dateext
# uncomment this if you want your log files compressed
#compress
# RPM packages drop log rotation information into this directory
include /etc/logrotate.d
# system-specific logs may be also be configured here.
```


With the above settings:

- Logs are backed up weekly.
- Data upto 4 weeks old is backed up. Any log backup created before that will be deleted.

**Caution:** Since these logs take up `/var` space, sufficient space must be allocated to `/var` partition if being created. If `/var` partition space fills up, control plane might crash.

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

## 4.2 Control plane logs

All log files can be viewed via CLI. Alternatively, most log files can be viewed via the Dashboard tab (  ) on the grafana UI.

**Caution:** It is not recommended to delete the below log files or the directories they reside in.

---


**Note:** Log files are rotated periodically as a storage consideration. To customize how often logs are rotated, edit the `/etc/logrotate.conf` file on the node.

---

Below is a list of all logs available to Loki and can be accessed on the dashboard:

Table 1: Log files

| Name                    | Location   | Purpose  | Additional Information   |
|-------------------------|--|--|--|
| Omnia Logs              | /var/log/omnia.log                               | Omnia Log  | This log is configured by Default. This log can be used to track all changes made by all playbooks in the omnia directory. |
| Accelerator Logs        | /var/log/omnia/accel                             | Accelerator Log                                  | This log is configured by Default.   |
| Monitor Logs            | /var/log/omnia/moni                              | Monitor Log                                      | This log is configured by Default.   |
| Network Logs            | /var/log/omnia/netwo                             | Network Log                                      | This log is configured by Default.   |
| Platform Logs           | /var/log/omnia/platfo                            | Platform Log                                     | This log is configured by Default.   |
| Provision Logs          | /var/log/omnia/provi                             | Provision Log                                    | This log is configured by Default.   |
| Scheduler Logs          | /var/log/omnia/sched                             | Scheduler Log                                    | This log is configured by Default.   |
| Security Logs           | /var/log/omnia/secur                             | Security Log                                     | This log is configured by Default.   |
| Storage Logs            | /var/log/omnia/stora                             | Storage Log                                      | This log is configured by Default.   |
| Telemetry Logs          | /var/log/omnia/telem                             | Telemetry Log                                    | This log is configured by Default.   |
| Utils Logs              | /var/log/omnia/utills.l                          | Utils Log  | This log is configured by Default.   |
| Cluster Utilities Logs  | /var/log/omnia/utills_                           | Cluster Utils Log                                | This log is configured by Default.   |
| syslogs                 | /var/log/messages                                | System Logging                                   | This log is configured by Default.   |
| Audit Logs              | /var/log/audit/audit.l                           | All Login Attempts                               | This log is configured by Default.   |
| CRON logs               | /var/log/cron                                    | CRON Job Logging                                 | This log is configured by Default.   |
| Pods logs               | /var/log/pods/ * / * / * log                     | k8s pods   | This log is configured by Default.   |
| Access Logs             | /var/log/dirsrv/slapd-<br><Realm<br>Name>/access | Directory Server Utilization                     | This log is available when FreeIPA or 389ds is set up ( ie when enable_security_support is set to 'true').                 |
| Error Log               | /var/log/dirsrv/slapd-<br><Realm<br>Name>/errors | Directory Server Errors                          | This log is available when FreeIPA or 389ds is set up ( ie when enable_security_support is set to 'true').                 |
| CA Transaction Log      | /var/log/pki/pki-tomcat/ca/transaction           | FreeIPA PKI Transactions                         | This log is available when FreeIPA or 389ds is set up ( ie when enable_security_support is set to 'true').                 |
| KRB5KDC                 | /var/log/krb5kdc.log                             | KDC Utilization                                  | This log is available when FreeIPA or 389ds is set up ( ie when enable_security_support is set to 'true').                 |
| Secure logs             | /var/log/secure                                  | Login Error Codes                                | This log is available when FreeIPA or 389ds is set up ( ie when enable_security_support is set to 'true').                 |
| HTTPD logs              | /var/log/httpd/ *                                | FreeIPA API Calls                                | This log is available when FreeIPA or 389ds is set up ( ie when enable_security_support is set to 'true').                 |
| DNF logs                | /var/log/dnf.log                                 | Installation Logs                                | This log is configured on Rocky OS.  |
| BeeGFS Logs             | /var/log/beegfs-client.log                       | BeeGFS Logs                                      | This log is configured on BeeGFS client nodes.   |
| Compute Logs            | /var/log/xcat/comput                             | Logs system messages from all cluster nodes.     | This log is configured by Default.   |
| Cluster deployment logs | /var/log/xcat/cluster.l                          | Logs deployment messages from all cluster nodes. | This log is configured by Default.   |



### 4.3 Logs of individual containers

1. A list of namespaces and their corresponding pods can be obtained using: `kubectl get pods -A`
2. Get a list of containers for the pod in question using: `kubectl get pods <pod_name> -o jsonpath='{.spec.containers[*].name}'`
3. Once you have the namespace, pod and container names, run the below command to get the required logs:  
`kubectl logs pod <pod_name> -n <namespace> -c <container_name>`

### 4.4 Provisioning logs

Logs pertaining to actions taken during `provision.yml` can be viewed in `/var/log/xcat/cluster.log` and `/var/log/xcat/computes.log` on the control plane.

**Note:** As long as a node has been added to a cluster by Omnia, deployment events taking place on the node will be updated in `/var/log/xcat/cluster.log`.

## 4.5 Telemetry logs

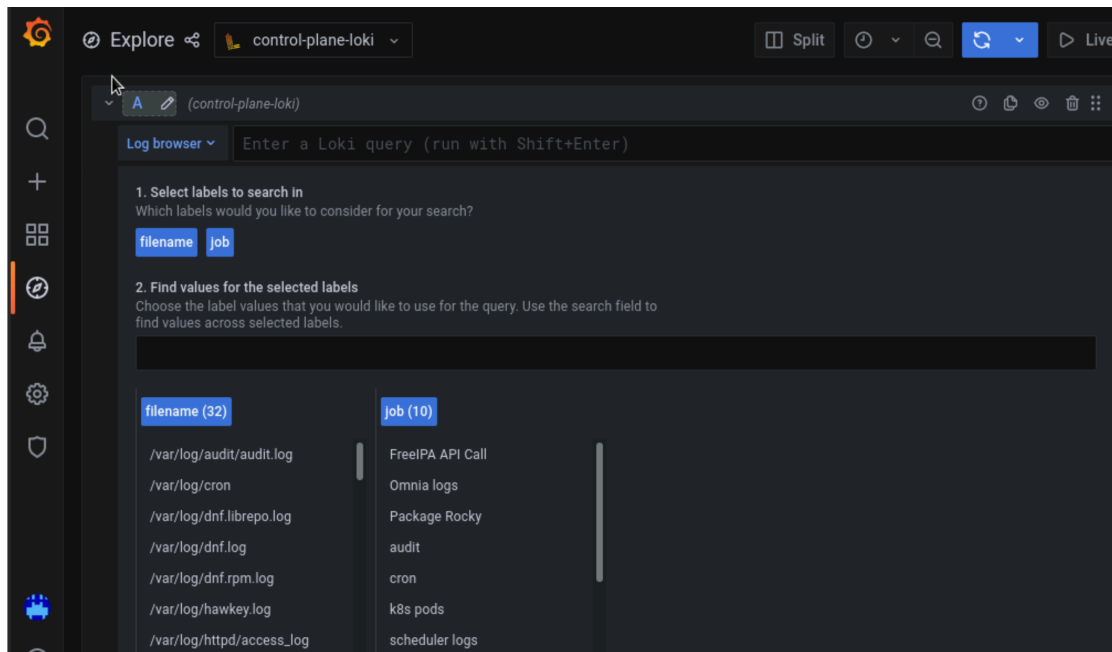
Logs pertaining to actions taken by Omnia or iDRAC telemetry can be viewed in `/var/log/messages`. Each log entry is tagged “omnia\_telemetry”. Log entries typically follow this format.

```
<Date time> <Node name> omnia_telemetry[<Process ID>]: <name of file>:<name of method_
→throwing error>: <Error message>
```

## 4.6 Grafana Loki

After `telemetry.yml` is run, Grafana services are installed on the control plane.

- i. Get the Grafana IP using `kubectl get svc -n grafana`.
- ii. Login to the Grafana UI by connecting to the cluster IP of grafana service via port 5000. That is `http://xx.xx.xx.xx:5000/login`.
- iii. In the Explore page, select **control-plane-loki**.



- iv. The log browser allows users to filter logs by job, node, user, etc. Ex:

```
(job= "cluster deployment logs") |= "nodename"
(job="compute log messages") |= "nodename" |= "node_username"
```

Custom dashboards can be created per user requirement.

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).



## TROUBLESHOOTING

## 5.1 Known issues

**Why doesn't my newly discovered server list a MAC ID in the cluster.nodeinfo table?**

Due to internal MAC ID conflicts on the target nodes, the MAC address will be listed against the target node using this format `MAC ADDRESS 1 | MAC ADDRESS 2! *NOIP*` in the xCAT node object.

```
[root@ ]# lsdef | grep mac
mac=f4:02:70:b8:cc:80|f4:02:70:f1:7e:a3!*NOIP*
```

**Hostname assignment by Omnia is not sequential (that is from 1-xx) when discovering nodes via BMC or switch-based methods**

Omnia does not maintain any order when assigning hostnames to target nodes.

**cluster nodes get updated to Rocky 8.8 automatically irrespective of the input parameters provided to provision\_config.yml when provision\_os is Rocky.**

**Potential Cause:** In Rocky Linux, online repos are enabled by default, and they always point to the latest Rocky repository (currently Rocky Linux 8.8).

**Resolution:** This will be addressed in a later release.

**Why does the task Assign admin NIC IP fail during provision.yml with errors?**

```
SK [provision_validation : Failed - Assign admin nic IP] *****
sk path: /root/shubh/omnia/provision/roles/provision_validation/tasks/assign_admin_nic_config.yml:25
tal: [localhost]: FAILED! => {"changed": false, "msg": "Failed. Please assign admin nic IP. "}

SK [provision_validation : Failed - Assign admin nic IP] *****
sk path: /root/shubh/omnia/provision/roles/provision_validation/tasks/assign_admin_nic_config.yml:37
tal: [localhost]: FAILED! => {"changed": false, "msg": "Failed. Please assign admin nic IP. "}

AY RECAP *****
calhost          : ok=33  changed=4  unreachable=0  failed=1  skipped=8  rescued=1  ignored=0
```

**Potential Cause:** Omnia validates the admin NIC IP on the control plane. If the user has not assigned an admin NIC IP in case of dedicated network interface type, an error message is returned. There is a parsing logic that is being applied on the blank IP and hence, the error displays twice.

**Resolution:** Ensure a control plane IP is assigned to the admin NIC.

**Kubernetes pods on the manager node are in CreateContainerConfigError and Calico Pods are in CrashLoopBackoff error after running omnia.yml.**

**Potential Cause:**

Calico pods are configured with the NIC name of the manager node. If the NIC name of the other nodes are not the same, the pods will throw an error and retry later.

**Workaround:**

The manager and cluster nodes should have connectivity over the same admin NIC.

**Why does the task ``cluster\_preperation : Install sshpass`` fail during ``omnia.yml`` on cluster nodes running RHEL 8.5 and below versions.**

**Potential Cause:**

- sshpass is not available in any of the repositories on the control plane.

**Resolution:**

- Enable RedHat subscription or ensure that sshpass is available to install or download to the control plane (from any local repository).

**Why are some target servers not reachable after running PXE booting them?**

**Potential Causes:**

1. The server hardware does not allow for auto rebooting
2. PXE booting is hung on the node

**Resolution:**

1. Login to the iDRAC console to check if the server is stuck in boot errors (F1 prompt message). If true, clear the hardware error or disable POST (PowerOn Self Test).
2. Hard-reboot the server to bring up the server and verify that the boot process runs smoothly. (If it gets stuck again, disable PXE and try provisioning the server via iDRAC.)

**Why does the task ‘Provision: Fetch the available subnets and netmasks’ fail with ‘no ipv4\_secondaries present’?**

```
ane_common : Fetch the available subnets and netmasks] *****
t]: FAILED! => {"msg": "The task includes an option with an undefined variable. The error was: 'dict ob
daries'\n\nThe error appears to be in '/root/omnia_mgmt_lom/omnia/control_plane/roles/control_plane_co
.yml': line 75, column 3, but may\nbe elsewhere in the file depending on the exact syntax problem.\n\nT
\n\n- name: Fetch the available subnets and netmasks\n  ^ here\n"}
```

**Potential Cause:** If a shared LOM environment is in use, the management network/host network NIC may only have one IP assigned to it.

**Resolution:** Ensure that the NIC used for host and data connections has 2 IPs assigned to it.

**Why does provisioning RHEL 8.3 fail on some nodes with “dbus.error.DBusError: ‘NoneType’ object has no attribute ‘set\_property’”?**

This error is known to RHEL and is being addressed [here](#). Red Hat has offered a user intervention [here](#). Omnia recommends that in the event of this failure, any OS other than RHEL 8.3.

**Why is the Infiniband NIC down after provisioning the server?**

For servers running Rocky, enable the Infiniband NIC manually, use `ifup <InfiniBand NIC>`.

Alternatively, run `network.yml` or `post_provision.yml` (Only if the nodes are provisioned using Omnia) to activate the NIC.

**Why does the Task [infiniband\_switch\_config : Authentication failure response] fail with the message ‘Status code was -1 and not [302]: Request failed: <urlopen error [Errno 111] Connection refused>’ on Infiniband Switches when running infiniband\_switch\_config.yml?**

To configure a new Infiniband Switch, it is required that HTTP and JSON gateway be enabled. To verify that they are enabled, run:

`show web` (To check if HTTP is enabled)



`show json-gw` (To check if JSON Gateway is enabled)

To correct the issue, run:

`web http enable` (To enable the HTTP gateway)

`json-gw enable` (To enable the JSON gateway)

#### Why does the task ‘Initialize kubeadm’ fail while running monitor.yml?

This issue is caused by incompatibility between Rocky 8.7 and kubernetes due to cri-o. For more information, [click here](#).

#### Why does PXE boot fail with tftp timeout or service timeout errors?

##### Potential Causes:

- RAID is configured on the server.
- Two or more servers in the same network have xCAT services running.
- The target cluster node does not have a configured PXE device with an active NIC.

##### Resolution:

1. Create a Non-RAID or virtual disk on the server.
2. Check if other systems except for the control plane have xcatd running. If yes, then stop the xCAT service using the following commands: `systemctl stop xcatd`.
3. On the server, go to BIOS Setup -> Network Settings -> PXE Device. For each listed device (typically 4), configure an active NIC under PXE device settings

#### Why do Kubernetes Pods show “ImagePullBack” or “ErrPullImage” errors in their status?

##### Potential Cause:

- The errors occur when the Docker pull limit is exceeded.

##### Resolution:

- For `omnia.yml` and `provision.yml` : Provide the docker username and password for the Docker Hub account in the `omnia_config.yml` file and execute the playbook.
- For HPC cluster, during `omnia.yml` execution, a kubernetes secret ‘dockerregcred’ will be created in default namespace and patched to service account. User needs to patch this secret in their respective namespace while deploying custom applications and use the secret as `imagePullSecrets` in yaml file to avoid `ErrImagePull`. [Click here for more info](<https://kubernetes.io/docs/tasks/configure-pod-container/pull-image-private-registry/>)

---

**Note:** If the playbook is already executed and the pods are in **ImagePullBack** state, then run `kubeadm reset -f` in all the nodes before re-executing the playbook with the docker credentials.

---

#### Why does the task ‘Gather facts from all the nodes’ stuck when re-running ``omnia.yml``?

**Potential Cause:** Corrupted entries in the `/root/.ansible/cp/` folder. For more information on this issue, [check this out!](#)

**Resolution:** Clear the directory `/root/.ansible/cp/` using the following commands:

```
cd /root/.ansible/cp/

rm -rf *
```

Alternatively, run the task manually:

```
cd omnia/utils/cluster
ansible-playbook gather_facts_resolution.yml
```

**What to do after a reboot if kubectl commands return: ``The connection to the server head\_node\_ip:port was refused - did you specify the right host or port?``**

On the control plane or the manager node, run the following commands:

```
swapoff -a

systemctl restart kubelet
```

**What to do if the nodes in a Kubernetes cluster reboot:**

Wait for 15 minutes after the Kubernetes cluster reboots. Next, verify the status of the cluster using the following commands:

- `kubectl get nodes` on the manager node to get the real-time k8s cluster status.
- `kubectl get pods all-namespaces` on the manager node to check which the pods are in the **Running** state.
- `kubectl cluster-info` on the manager node to verify that both the k8s master and kubeDNS are in the **Running** state.

**What to do when the Kubernetes services are not in the Running state:**

1. Run `kubectl get pods all-namespaces` to verify that all pods are in the **Running** state.
2. If the pods are not in the **Running** state, delete the pods using the command: `kubectl delete pods <name of pod>`
3. Run the corresponding playbook that was used to install Kubernetes: `omnia.yml`, `jupyterhub.yml`, or `kubeflow.yml`.

**Why do Kubernetes Pods stop communicating with the servers when the DNS servers are not responding?**

**Potential Cause:** The host network is faulty causing DNS to be unresponsive

**Resolution:**

1. In your Kubernetes cluster, run `kubeadm reset -f` on all the nodes.
2. On the management node, edit the `omnia_config.yml` file to change the Kubernetes Pod Network CIDR. The suggested IP range is 192.168.0.0/16. Ensure that the IP provided is not in use on your host network.
3. Set `scheduler_type: "k8s"` in `input/omnia_config.yml` and run `omnia.yml`.

**Why does pulling images to create the Kubeflow timeout causing the ‘Apply Kubeflow Configuration’ task to fail?**

**Potential Cause:** Unstable or slow Internet connectivity.

**Resolution:**

1. Complete the PXE booting/format the OS on the manager and cluster nodes.
2. In the `omnia_config.yml` file, change the `k8s_cni` variable value from `calico` to `flannel`.
3. Run the Kubernetes and Kubeflow playbooks.

**Why does the ‘Initialize Kubeadm’ task fail with ‘nnode.Registration.name: Invalid value: "<Host name>"’?**

**Potential Cause:** The `control_plane` playbook does not support hostnames with an underscore in it such as `‘mgmt_station’`.

As defined in RFC 822, the only legal characters are the following: 1. Alphanumeric (a-z and 0-9): Both uppercase and lowercase letters are acceptable, and the hostname is case-insensitive. In other words, omnia.test is identical to OMNIA.TEST and Omnia.test.

2. Hyphen (-): Neither the first nor the last character in a hostname field should be a hyphen.
3. Period (.): The period should be used only to delimit fields in a hostname (e.g., dvader.empire.gov)

**What to do when Kubeflow pods are in ‘ImagePullBackOff’ or ‘ErrImagePull’ status after executing kube-flow.yml:**

**Potential Cause:** Your Docker pull limit has been exceeded. For more information, click [here](https://www.docker.com/increase-rate-limits)

1. Delete Kubeflow deployment by executing the following command in manager node: `kfctl delete -V -f /root/k8s/omnia-kubeflow/kfctl_k8s_istio.v1.0.2.yaml`
2. Re-execute kube-flow.yml after 8-9 hours

**What to do when omnia.yml fail with ‘Error: kinit: Connection refused while getting default ccache’ while completing the security role?**

1. Start the sssd-kcm.socket: `systemctl start sssd-kcm.socket`
2. Re-run omnia.yml

**What to do when Slurm services do not start automatically after the cluster reboots:**

- Manually restart the slurmd services on the manager node by running the following commands:

```
systemctl restart slurmdbd
systemctl restart slurmctld
systemctl restart prometheus-slurm-exporter
```

- Run `systemctl status slurmd` to manually restart the following service on all the cluster nodes.

**Why do Slurm services fail?**

**Potential Cause:** The `slurm.conf` is not configured properly.

Recommended Actions:

1. Run the following commands:

```
slurmdbd -Dvvv
slurmctld -Dvvv
```

2. Refer the `/var/lib/log/slurmctld.log` file for more information.

**What causes the “Ports are Unavailable” error?**

**Potential Cause:** Slurm database connection fails.

Recommended Actions:

1. Run the following commands::

```
slurmdbd -Dvvv
slurmctld -Dvvv
```

2. Refer the `/var/lib/log/slurmctld.log` file.
3. Check the output of `netstat -antp | grep LISTEN` for PIDs in the listening state.
4. If PIDs are in the **Listening** state, kill the processes of that specific port.

5. Restart all Slurm services:

```
slurmctl restart slurmctld on manager node

systemctl restart slurmdbd on manager node

systemctl restart slurmd on compute node
```

**Why does the task ‘nfs\_client: Mount NFS client’ fail with ‘Failed to mount NFS client. Make sure NFS Server is running on IP xx.xx.xx.xx’?**

**Potential Cause:**

- The required services for NFS may not be running:
  - nfs
  - rpc-bind
  - mountd

**Resolution:**

- Enable the required services using `firewall-cmd --permanent --add-service=<service name>` and then reload the firewall using `firewall-cmd --reload`.

**What to do when omnia.yml fails with nfs-server.service might not be running on NFS Server. Please check or start services`?**

**Potential Cause:** nfs-server.service is not running on the target node.

**Resolution:** Use the following commands to bring up the service:

```
systemctl start nfs-server.service

systemctl enable nfs-server.service
```

**Why does the task ‘Install Packages’ fail on the NFS node with the message: ‘Failure in talking to yum: Cannot find a valid baseurl for repo: base/7/x86\_64.’**

**Potential Cause:**

There are connections missing on the NFS node.

**Resolution:**

Ensure that there are 3 NICs being used on the NFS node:

1. For provisioning the OS
2. For connecting to the internet (Management purposes)
3. For connecting to PowerVault (Data Connection)

**Why do pods and images appear to get deleted automatically?**

**Potential Cause:**

Lack of space in the root partition (/) causes Linux to clear files automatically (Use `df -h` to diagnose the issue).

**Resolution:**

- Delete large, unused files to clear the root partition (Use the command `find / -xdev -size +5M | xargs ls -lh | sort -n -k5` to identify these files). Before running `monitor.yml`, it is recommended to have a minimum of 50% free space in the root partition.

- Once the partition is cleared, run `kubeadm reset -f`
- Re-run `monitor.yml`

**What to do when the JupyterHub or Prometheus UI is not accessible:**

Run the command `kubect1 get pods namespace default` to ensure **nfs-client** pod and all Prometheus server pods are in the **Running** state.

**What to do if PowerVault throws the error: ``Error: The specified disk is not available. - Unavailable disk (0.x) in disk range '0.x-x'``:**

1. Verify that the disk in question is not part of any pool: `show disks`
2. If the disk is part of a pool, remove it and try again.

**Why does PowerVault throw the error: ``You cannot create a linear disk group when a virtual disk group exists on the system.``?**

At any given time only one type of disk group can be created on the system. That is, all disk groups on the system have to exclusively be linear or virtual. To fix the issue, either delete the existing disk group or change the type of pool you are creating.

**Why does the task 'nfs\_client: Mount NFS client' fail with ``No route to host``?****Potential Cause:**

- There's a mismatch in the share path listed in `/etc/exports` and in `omnia_config.yml` under `nfs_client_params`.

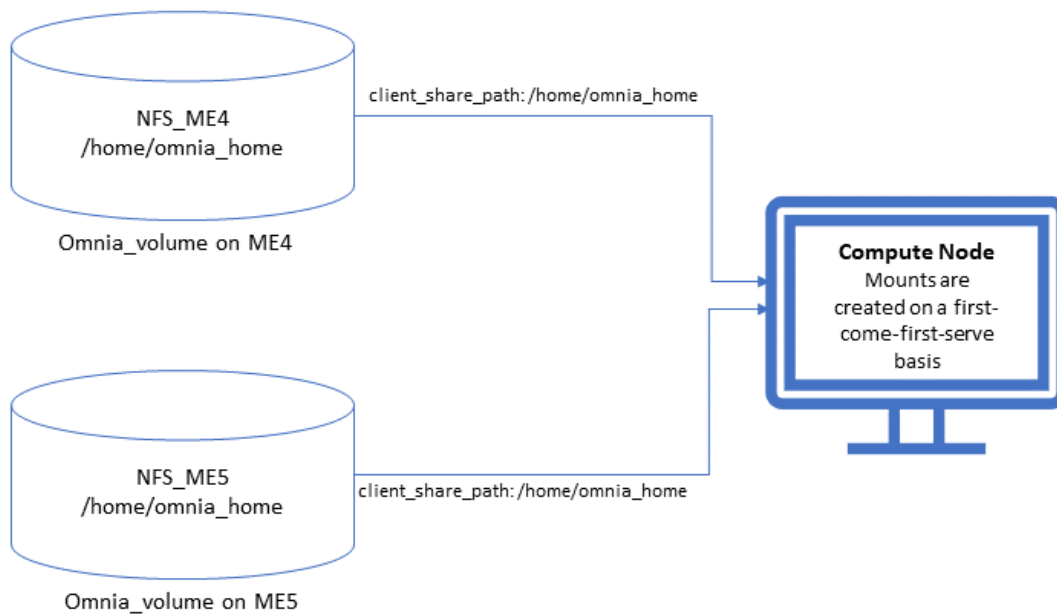
**Resolution:**

- Ensure that the input paths are a perfect match down to the character to avoid any errors.

**Why is my NFS mount not visible on the client?**

**Potential Cause:** The directory being used by the client as a mount point is already in use by a different NFS export.

**Resolution:** Verify that the directory being used as a mount point is empty by using `cd <client share path> | ls` or `mount | grep <client share path>`. If empty, re-run the playbook.



### Why does the ``BeeGFS-client`` service fail?

#### Potential Causes:

1. SELINUX may be enabled. (use `sestatus` to diagnose the issue)
2. Ports 8008, 8003, 8004, 8005 and 8006 may be closed. (use `systemctl status beegfs-mgmt`, `systemctl status beegfs-meta`, `systemctl status beegfs-storage` to diagnose the issue)
3. The BeeGFS set up may be incompatible with RHEL.

#### Resolution:

1. If SELinux is enabled, update the file `/etc/sysconfig/selinux` and reboot the server.
2. Open all ports required by BeeGFS: 8008, 8003, 8004, 8005 and 8006
3. Check the [support matrix for RHEL or Rocky](../Support\_Matrix/Software/Operating\_Systems) to verify your set-up.
4. For further insight into the issue, check out `/var/log/beegfs-client.log` on nodes where the BeeGFS client is running.

### Why does the task 'security: Authenticate as admin' fail?

**Potential Cause:** The required services are not running on the node. Verify the service status using::

```
systemctl status sssd-kcm.socket
systemctl status sssd.service
```

#### Resolution:

- Restart the services using::

```
systemctl start sssd-kcm.socket
systemctl start sssd.service
```

- Re-run omnia.yml using:

```
ansible-playbook omnia.yml
```

### Why does installing FreeIPA fail on RHEL servers?

```
TASK [control_plane_security : Install ipa server in CentOS or Rocky] *****
task path: /root/omnia/control_plane/roles/control_plane_security/tasks/install_ipa_server.yml:17
fatal: [localhost]: FAILED! => {"censored": "the output has been hidden due to the fact that 'no_log: true' was
specified for this result", "changed": true}

TASK [control_plane_security : Install ipa server failed] *****
task path: /root/omnia/control_plane/roles/control_plane_security/tasks/install_ipa_server.yml:25
fatal: [localhost]: FAILED! => {"changed": false, "msg": "Error: ['Checking DNS domain omnia.test., please wait
...', 'Validate installation settings ...', 'Create file system structures ...', 'Perform SELinux labeling ...',
'Warning no default label for /dev/shm/slapd-OMNIA-TEST', 'Create database backend: dc=omnia,dc=test ...', 'Per
form post-installation tasks ...', 'CalledProcessError(Command ['\\bin/systemctl\\', '\\start\\', '\\pki-tomcat
d@pki-tomcat.service\\'] returned non-zero exit status 1: '\\Job for pki-tomcatd@pki-tomcat.service failed becau
se a timeout was exceeded.\\n\\nSee '\\systemctl status pki-tomcatd@pki-tomcat.service\\' and '\\journalctl -xe\\' fo
r details.\\n\\n\\')", 'The ipa-server-install command failed. See /var/log/ipaserver-install.log for more inform
ation']"}

PLAY RECAP *****
localhost : ok=312  changed=57  unreachable=0  failed=1  skipped=283  rescued=1  ignored
=0
```

**Potential Causes:** Required repositories may not be enabled by your red hat subscription.

**Resolution:** Enable all required repositories via your red hat subscription.

### Why would FreeIPA server/client installation fail?

**Potential Cause:**

The hostnames of the manager and login nodes are not set in the correct format.

**Resolution:**

If you have enabled the option to install the login node in the cluster, set the hostnames of the nodes in the format: *hostname.domainname*. For example, *manager.omnia.test* is a valid hostname for the login node. **Note:** To find the cause for the failure of the FreeIPA server and client installation, see *ipaserver-install.log* in the manager node or */var/log/ipaclient-install.log* in the login node.

### Why does FreeIPA installation fail on the control plane when the public NIC provided is static?

**Potential Cause:** The network config file for the public NIC on the control plane does not define any DNS entries.

**Resolution:** Ensure the fields DNS1 and DNS2 are updated appropriately in the file */etc/sysconfig/network-scripts/ifcfg-<NIC name>*.

### What to do when JupyterHub pods are in ‘ImagePullBackOff’ or ‘ErrImagePull’ status after executing jupyterhub.yml:

**Potential Cause:** Your Docker pull limit has been exceeded. For more information, [click here](#).

1. Delete Jupyterhub deployment by executing the following command in manager node: `helm delete jupyterhub -n jupyterhub`
2. Re-execute `jupyterhub.yml` after 8-9 hours.

### What to do if NFS clients are unable to access the share after an NFS server reboot?

Reboot the NFS server (external to the cluster) to bring up the services again:

```
systemctl disable nfs-server
systemctl enable nfs-server
systemctl restart nfs-server
```

**Why do Kuberneteschildnode & kubernetesnodes log as Pass in the database even if there are nodes in the Ready,Schedulingdisabled state?**

**Potential Cause:** Omnia telemetry considers Ready, SchedulingDisabled as a Ready state of Kubernetes nodes . So, even if the `kubectl get nodes` command shows any node's state as Ready, SchedulingDisabled, the entry in DB for Kuberneteschildnode & kubernetesnodes will be logged as Pass instead of Fail.

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

## 5.2 Frequently asked questions

**Why is the provisioning status of the target servers stuck at 'installing' in cluster.nodeinfo (omniadb)?**

```
omniadb=# select id,node,hostname,admin_ip,status from cluster.nodeinfo;
 id |      node      |      hostname      |      admin_ip      |      status
-----+-----+-----+-----+-----
  2 | node1          | node1.orchid.cluster | 10.5.0.1           |
  6 | control_plane | orchidcp.orchid.cluster | 10.5.255.254       |
  3 | node2          | node2.orchid.cluster | 10.5.0.2           | installing
  4 | node3          | node3.orchid.cluster | 10.5.0.3           | installing
  5 | node4          | node4.orchid.cluster | 10.5.0.4           | installing
(5 rows)
```

```
* When reporting a bug add logs from /tmp as separate text/plain attachments
21:59:02 Running pre-installation scripts
21:59:08 Not asking for UNC because of an automated install
21:59:08 Not asking for UNC because text mode was explicitly asked for in kickstart
Starting automated install...Saving storage configuration...
Checking storage configuration...
.

=====
=====
Installation

1) [x] Language settings                2) [x] Time settings
   (English (United States))            (CST6CDT timezone)
3) [!] Installation source              4) [!] Software selection
   (Error setting up software           (Error checking software
   source)                             selection)
5) [x] Installation Destination         6) [x] Kdump
   (Custom partitioning selected)       (Kdump is enabled)
7) [x] Network configuration           8) [ ] User creation
   (Wired (eno3) connected)             (No user will be created)

Please make a selection from the above ['b' to begin installation, 'q' to quit,
'r' to refresh]:
```

**Potential Causes:**

- Disk partition may not have enough storage space per the requirements specified in input/provision\_config (under disk\_partition)
- The provided ISO may be corrupt/incomplete.
- Hardware issues (Auto reboot may fail at POST)
- A virtual disk may not be created

**Resolution:**



- Add more space to the server or modify the requirements specified in `input/provision_config` (under `disk_partition`)
- Download the ISO again, verify the checksum/ download size and re-run the provision tool.
- Resolve/replace the faulty hardware and PXE boot the node.
- Create a virtual disk and PXE boot the node.

**Why is the provisioning status of my target servers stuck at ‘powering-on’ in the cluster.info (omniadb)?**

**Potential Cause:**

- Hardware issues (Auto-reboot may fail due to hardware tests failing)
- The target node may already have an OS and the first boot PXE device is not configured correctly.

**Resolution:**

- Resolve/replace the faulty hardware and PXE boot the node.
- Target servers should be configured to boot in PXE mode with the appropriate NIC as the first boot device.

**What to do if PXE boot fails while discovering target nodes via switch\_based discovery with provisioning status stuck at ‘powering-on’ in cluster.nodeinfo (omniadb):**

```
xcat.genesis.dodiscovery: Beginning echo information to discovery packet file...
xcat.genesis.dodiscovery: Discovery packet file is ready.
xcat.genesis.dodiscovery: Sending the discovery packet to xCAT (172.59.255.254:3001)...
xcat.genesis.dodiscovery: Sleeping 5 seconds...
xcat.genesis.minixcatd: The request is processing by xCAT master...
xcat.genesis.minixcatd: The request is already processed by xCAT master, but not matched.
xcat.genesis.dodiscovery: Beginning echo information to discovery packet file...
xcat.genesis.dodiscovery: Discovery packet file is ready.
xcat.genesis.dodiscovery: Sending the discovery packet to xCAT (172.59.255.254:3001)...
xcat.genesis.dodiscovery: Sleeping 5 seconds...
xcat.genesis.minixcatd: The request is processing by xCAT master...
xcat.genesis.minixcatd: The request is already processed by xCAT master, but not matched.
xcat.genesis.dodiscovery: Beginning echo information to discovery packet file...
xcat.genesis.dodiscovery: Discovery packet file is ready.
xcat.genesis.dodiscovery: Sending the discovery packet to xCAT (172.59.255.254:3001)...
xcat.genesis.dodiscovery: Sleeping 5 seconds...
xcat.genesis.minixcatd: The request is processing by xCAT master...
xcat.genesis.minixcatd: The request is already processed by xCAT master, but not matched.
xcat.genesis.dodiscovery: Beginning echo information to discovery packet file...
xcat.genesis.dodiscovery: Discovery packet file is ready.
xcat.genesis.dodiscovery: Sending the discovery packet to xCAT (172.59.255.254:3001)...
xcat.genesis.dodiscovery: Sleeping 5 seconds...
xcat.genesis.minixcatd: The request is processing by xCAT master...
xcat.genesis.minixcatd: The request is already processed by xCAT master, but not matched.
xcat.genesis.dodiscovery: Beginning echo information to discovery packet file...
xcat.genesis.dodiscovery: Discovery packet file is ready.
```

1. Rectify any probable causes like incorrect/unavailable credentials (`switch_snmp3_username` and `switch_snmp3_password` provided in `input/provision_config.yml`), network glitches or incorrect switch IP/port details.
2. Run the clean up script by:

```
cd utils
ansible-playbook control_plane_cleanup.yml
```

3. Re-run the provision tool (`ansible-playbook provision.yml`).

**What to do if playbook execution fails due to external (network, hardware etc) failure:**

Re-run the playbook whose execution failed once the issue is resolved.

**Why don't IPA commands work after setting up FreeIPA on the cluster?**

**Potential Cause:**

Kerberos authentication may be missing on the target node.

**Resolution:**

Run `kinit admin` on the node and provide the `kerberos_admin_password` when prompted. (This password is also entered in `input/security_config.yml`.)

**Why are the status and admin\_mac fields not populated for specific target nodes in the cluster.nodeinfo table?****Causes:**

- Nodes do not have their first PXE device set as designated active NIC for PXE booting.
- Nodes that have been discovered via SNMP or mapping file have not been PXE booted.

**Resolution:**

- Configure the first PXE device to be active for PXE booting.
- PXE boot the target node manually.

**Why does the provision.yml fail at 'provision validation: Install common packages for provision' on RHEL nodes running 8.5 or earlier?**

```
TASK [provision_validation : Install common packages for provision] *****
task path: /root/omnia/provision/roles/provision_validation/tasks/package_installation.yml:20
fatal: [localhost]: FAILED! => {"changed": false, "failures": ["No package sshpass available."], "msg": "Failed to install some of the specified packages", "rc": 1, "results": []}

PLAY RECAP *****
localhost                : ok=4    changed=0    unreachable=0    failed=1    skipped=0    rescued=0    ignored=0
```

**Potential Cause:**

- sshpass is not available in any of the repositories on the control plane.

**Resolution:**

- Enable RedHat subscription or ensure that sshpass is available to install or download to the control plane (from any local repository).

---

**Note:** This error can also take place when task `cluster_preparation : Install sshpass` is executed during `omnia.yml`.

---

What to do if user login fails when accessing a cluster node:

```
[root@springcp ~]# ssh user1@10.27.0.2
@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@
@  WARNING: REMOTE HOST IDENTIFICATION HAS CHANGED!  @
@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@
IT IS POSSIBLE THAT SOMEONE IS DOING SOMETHING NASTY!
Someone could be eavesdropping on you right now (man-in-the-middle attack)!
It is also possible that a host key has just been changed.
The fingerprint for the ECDSA key sent by the remote host is
SHA256:zJq2bCTNDR+rR3zi1Kw3kHxYyPXZI50EyEvByGwh6Eg.
Please contact your system administrator.
Add correct host key in /root/.ssh/known_hosts to get rid of this message.
Offending ECDSA key in /root/.ssh/known_hosts:20
Password authentication is disabled to avoid man-in-the-middle attacks.
Keyboard-interactive authentication is disabled to avoid man-in-the-middle attacks.
user1@10.27.0.2: Permission denied (publickey,gssapi-keyex,gssapi-with-mic,password).
```

**Potential Cause:**

- ssh key on the control plane may be outdated.

**Resolution:**

- Refresh the key using `ssh-keygen -R <hostname/server IP>`.

- Retry login.

### Why does the ‘Fail if LDAP home directory exists’ task fail during user\_passwordless\_ssh.yml?

```
TASK [passwordless_ssh_ldap : Initialize username] *****
ok: [10.5.0.7]

TASK [passwordless_ssh_ldap : Check ldap home directory exists - user1] *****
ok: [10.5.0.7]

TASK [passwordless_ssh_ldap : Fail if ldap home directory not exists - user1] *****
fatal: [10.5.0.7]: FAILED! => {"changed": false, "msg": "Failed. ldap home directory /home/omnia-share//user1 not present for the user user1 in ldap server. Login to one of the nodes other than ldap server as the user so that home directory is created. Make sure ldap home directory /home/omnia-share/ mounted as nfs share in ldap server."}

PLAY RECAP *****
10.5.0.1      : ok=22  changed=4    unreachable=0    failed=0    skipped=15   rescued=0     ignored=0
10.5.0.2      : ok=14  changed=4    unreachable=0    failed=0    skipped=2    rescued=0     ignored=0
10.5.0.7      : ok=8   changed=0    unreachable=0    failed=1    skipped=1    rescued=0     ignored=0
localhost    : ok=16  changed=2    unreachable=0    failed=0    skipped=6    rescued=0     ignored=0
```

**Potential Cause:** The required NFS share is not set up on the control plane.

**Resolution:**

If `enable_omnia_nfs` is true in `input/omnia_config.yml`, follow the below steps to configure an NFS share on your LDAP server:

- **From the manager node:**

1. Add the LDAP server IP address to `/etc/exports`.
2. Run `exportfs -ra` to enable the NFS configuration.

- **From the LDAP server:**

1. Add the required `fstab` entries in `/etc/fstab` (The corresponding entry will be available on the compute nodes in `/etc/fstab`)
2. Mount the NFS share using `mount manager_ip: /home/omnia-share /home/omnia-share`

### Why does the ‘Import SCP from a local path’ task fail during idrac.yml?

```
TASK [provision_idrac : Import SCP from a local path and wait for this job to get completed] ***
FAILED - RETRYING: [165.29.0.106]: Import SCP from a local path and wait for this job to get completed (5 retries left).
FAILED - RETRYING: [165.29.0.106]: Import SCP from a local path and wait for this job to get completed (4 retries left).
FAILED - RETRYING: [165.29.0.106]: Import SCP from a local path and wait for this job to get completed (3 retries left).
FAILED - RETRYING: [165.29.0.106]: Import SCP from a local path and wait for this job to get completed (2 retries left).
FAILED - RETRYING: [165.29.0.106]: Import SCP from a local path and wait for this job to get completed (1 retries left).
fatal: [165.29.0.106]: FAILED! => {"attempts": 5, "changed": false, "msg": "Failed to import scp.", "scp_status": {"Data": {"Message": "returned status code doesn't match with the expected success code", "Status": "Failed", "StatusCode": 503}, "Message": "none", "Status": "Failed", "StatusCode": 503, "error": {"error": [{"@Message.ExtendedInfo": [{"Message": "A job operation is already running. Retry the operation after the existing job is completed.", "MessageArgs": [], "MessageArgs@odata.count": 0, "MessageId": "IDRAC.2.3.RAC0679", "RelatedProperties": [], "RelatedProperties@odata.count": 0, "Resolution": "Wait until the running job is completed or delete the scheduled job and retry the operation.", "Severity": "Warning"}]}, "code": "Base.1.7.GeneralError", "message": "A general error has occurred. See ExtendedInfo for more information"}}, "file": "/root/omnia/control-plane/roles/provision_idrac/files/idrac_scp.yml" "retval": true}}
```

**Potential Cause:** The target server may be hung during the booting process.

**Resolution:** Bring the target node up and re-run the script.

### Why does the ‘Verify primary\_dns is reachable’ task fail during provision.yml?

```
TASK [provision_validation : Verify primary_dns is reachable] *****
task path: /root/omnia/provision/roles/provision_validation/tasks/validate_provision_vars.yml:133
fatal: [localhost]: FAILED! => {"changed": false, "msg": "Failed. primary_dns is not reachable. Provide valid dns"}

PLAY RECAP *****
localhost    : ok=89  changed=7    unreachable=0    failed=1    skipped=26   rescued=0     ignored=0
```

Currently, the `primary_dns` value stored in `input/provision_config.yml` cannot be part of any of the subnets (`admin_nic_subnet`, `ib_nic_subnet` and `bmc_nic_subnet`) also defined in `input/provision_config.yml`.

Ex: If the `primary_dns` is set to 10.15.0.7, the subnet 10.15.0.0 cannot be used for `admin_nic_subnet`, `ib_nic_subnet` or `bmc_nic_subnet`.

### Why is the node status stuck at 'powering-on' or 'powering-off' after a control plane reboot?

**Potential Cause:** The nodes were powering off or powering on during the control plane reboot/shutdown.

**Resolution:** In the case of a planned shutdown, ensure that the control plane is shut down after the compute nodes. When powering back up, the control plane should be powered on and xCAT services resumed before bringing up the compute nodes. In short, have the control plane as the first node up and the last node down.

For more information, [click here](#)

### Why do subscription errors occur on RHEL control planes when `rhel_repo_local_path` (in `input/provision config.yml`) is not provided & control plane does not have an active subscription?

```
TASK [provision_validation : Subscription is not enabled] *****
task path: /root/omnia/provisioning/roles/provision_validation/tasks/validate_repo_path.yml:79
fatal: [localhost]: FAILED! => {"changed": false, "msg": "Failed. RedHat subscription not active. Activate RedHat
subscription or provide repos details in rhel repo path variable in provision config.yml"}
```

For many of Omnia's features to work, RHEL control planes need access to the following repositories:

1. AppStream
2. BaseOS
3. CRB

This can only be achieved using local repos specified in `rhel_repo_local_path` (`input/provision_config.yml`) OR having an active, available RedHat subscription.

**Note:** To enable the repositories, run the following commands:

```
subscription-manager repos --enable=codeready-builder-for-rhel-8-x86_64-rpms
subscription-manager repos --enable=rhel-8-for-x86_64-appstream-rpms
subscription-manager repos --enable=rhel-8-for-x86_64-baseos-rpms
```

Verify your changes by running:

```
yum repolist enabled
```

### Why does the task: Initiate reposync of AppStream, BaseOS and CRB fail?

[illegible]

**Potential Cause:** The `repo_url`, `repo_name` or `repo` provided in `rhel_repo_local_path` (`input/provision_config.yml`) may not be valid.

Omnia does not validate the input of `rhel_repo_local_path`.

**Resolution:** Ensure the correct values are passed before re-running `provision.yml`.

## How to add a new node for provisioning

1. Using a mapping file:

- Update the existing mapping file by appending the new entry (without the disrupting the older entries) or provide a new mapping file by pointing `pxe_mapping_file_path` in `provision_config.yml` to the new location.
- Run `provision.yml`.

2. Using the switch IP:

- Run `provision.yml` once the switch has discovered the potential new node.

### Why does the task: ‘BeeGFS: Rebuilding BeeGFS client module’ fail?

```
TASK [beegfs : Rebuilding BeegFS client module] *****task=beegfs/tasks/beegfs_client.yml:115
task path: /root/.omniqa/storage/roles/beegfs/tasks/beegfs_client.yml:115
fatal: [10.27.0.1]: FAILED! => {changed: true, cmd": ["etcd/init.d/beegfs-client", "rebuild"], "delta": "00:00:06.002034", "end": "2023-04-13 23:31:16.418490
", "msg": "non-zero return code", "rc": 2, "start": "2023-04-13 23:31:10.416370"}, stderr": "In file included from /opt/beegfs/src/client/client_module_7/bui
ld/../source/common/toolkit/Serialization.h:27,\n                                from /opt/beegfs/src/client/client_module_7/build/../source/common/net/message/control../Ne
tMessage.h:6,\n                                from /opt/beegfs/src/client/client_module_7/build/../source/common/net/message/control../SimpleStringMsg.h:4,\n                                from /opt/beegfs/src/client/client_module_7/build/../source/common/net/message/control/AckMsgEx.h:4,\n                                from /opt/beegfs/src/client/client_module_7/build/../source/os/OsCompat.h:11:10: fatal error: asm/kmap_types.h:
No such file or directory\n#include <asm/kmap_types.h>\n~~~~~\ncompilation terminated.\nmake[3]: *** [/opt/beegfs/src/client/client_module_7/build/../source/net/filesystem/FhgfsOpsCommKit.o] Error 1\nIn file included from /opt/beegfs/src/client/client_module_7/build/../source/net/filesystem/FhgfsOpsCommKit.o] Error 1\nIn file included from /opt/beegfs/src/client/client_module_7/build/../source/common/toolkit/Serialization.h:27,\n                                from /opt/beegfs/src/client/client_module_7/build/../source/common/net/message/session/FsYncLocalFileMsg.h:4,\n                                from /opt/beegfs/src/client/client_module_7/build/../source/net/filesystem/FhgfsOpsCommKit.c:2:\n/opt/beegfs/src/client/client_module_7/build/../source/os/OsCompat.h:11:10: fatal error: asm/kmap_types.h: No such file or directory\n#include <asm/kmap_types.h>\n~~~~~\ncompilation terminated.\nmake[3]: *** [/opt/beegfs/src/client/client_module_7/build/../source/net/filesystem/FhgfsOpsCommKit.o] Error 1\nIn file included from /opt/beegfs/src/client/client_module_7/build/../source/common/net/message/NetMessage.h:6,\n                                from /opt/beegfs/src/client/client_module_7/build/../source/common/toolkit/LookupIntentPntMsg.h:4,\n                                from /opt/beegfs/src/client/client_module_7/build/../source/common/toolkit/LookupIntentInfout.h:6,\n                                from /opt/beegfs/src/client/client_module_7/build/../source/filesystem/FhgfsOpShelper.h:6,\n                                from /opt/beegfs/src/client/client_module_7/build/../source/net/filesystem/FhgfsOpsRemoting.c:4:\n/opt/beegfs/src/client/client_module_7/build/../source/os/OsCompat.h:11:10: fatal error: asm/kmap_types.h: No such file or directory\n#include <asm/kmap_types.h>\n~~~~~\ncompilation terminated.\nmake[3]: *** [/opt/beegfs/src/client/client_module_7/build/../source/net/filesystem/FhgfsOpsRemoting.o] Error 1\nIn file included from /opt/beegfs/src/client/client_module_7/b
uild/../source/common/toolkit/Serialization.h:27,\n                                from /opt/beegfs/src/client/client_module_7/build/../source/common/net/message/session/rw/ReadLocalFileV2Msg.h:4,\n                                from /m/opt/beegfs/src/client/client_module_7/build/../source/net/filesystem/FhgfsOpsCommKitVec.c:2:\n/opt/beegfs/src/client/client_module_7/build/../source/os/OsCompat.h:11:10: fatal error: asm/kmap_types.h: No such file or directory\n#include <asm/kmap_types.h>\n~~~~~\ncompilation terminated.\nmake[3]: *** [/opt/beegfs/src/client/client_module_7/build/../source/net/filesystem/FhgfsOpsCommKitVec.o] Error 1\nmake[2]: *** [M
akefile:1585: module_opt/beegfs/src/client/client_module_7/build/../source] Error 2\nmake[1]: *** [Makefile:161: module] Error 2\nmake: *** [AutoRebuild.mk:
34: auto_rebuild] Error 2", "stderr_lines": ["In file included from /opt/beegfs/src/client/client_module_7/build/../source/common/toolkit/Serialization.h:27,"
, "                                from /opt/beegfs/src/client/client_module_7/build/../source/common/net/message/control../NetMessage.h:6,", "                                from /opt/be
egfs/src/client/client_module_7/build/../source/common/net/message/control../SimpleStringMsg.h:4,", "                                from /opt/beegfs/src/client/client_module_7/b
uild/../source/common/net/message/control/AckMsgEx.h:4,", "                                from /opt/beegfs/src/client/client_module_7/build/../source/net/message/NetH
essageFactory.c:2:", "/opt/beegfs/src/client/client_module_7/build/../source/os/OsCompat.h:11:10: fatal error: asm/kmap_types.h: No such file or directory", "
#include <asm/kmap_types.h>", "                                compilation terminated.", "make[3]: *** [/opt/beegfs/src/client/cl
ient_module_7/build/../source/net/message/NetMessageFactory.o] Error 1", "make[3]: *** Waiting for unfinished jobs....", "In file included from /opt/beegfs/sr
c/client/client_module_7/build/../source/common/toolkit/Serialization.h:27,", "                                from /opt/beegfs/src/client/client_module_7/build/../source/co
mmon/net/message/NetMessage.h:6,", "                                from /opt/beegfs/src/client/client_module_7/build/../source/net/filesystem/FsYncLocalFileMsg
.h:4,", "                                from /opt/beegfs/src/client/client_module_7/build/../source/net/filesystem/FhgfsOpsCommKit.c:2:", "/opt/beegfs/src/client/client_modu
le_7/build/../source/os/OsCompat.h:11:10: fatal error: asm/kmap_types.h: No such file or directory", "include <asm/kmap_types.h>", "                                compilation terminated.", "make[3]: *** [/opt/beegfs/src/client/client_module_7/build/../source/net/filesystem/FhgfsOpsC
ommKit.o] Error 1", "In file included from /opt/beegfs/src/client/client_module_7/build/../source/common/toolkit/Serialization.h:27,", "                                from /o
```

**Potential Cause:** BeeGFS version 7.3.0 is in use.

**Resolution:** Use BeeGFS client version 7.3.1 when setting up BeeGFS on the cluster.

### Why does splitting an ethernet Z series port fail with “Failed. Either port already split with different breakout value or port is not available on ethernet switch”?

**Potential Cause:**

1. The port is already split.
2. It is an even-numbered port.

**Resolution:**

Changing the `breakout_value` on a split port is currently not supported. Ensure the port is un-split before assigning a new `breakout_value`.

### How to enable DHCP routing on Compute Nodes:

To enable routing, update the `primary_dns` and `secondary_dns` in `provision_config.yml` with the appropriate IPs (hostnames are currently not supported). For compute nodes that are not directly connected to the internet (ie only host network is configured), this configuration allows for internet connectivity.

### What to do if the LC is not ready:

- Verify that the LC is in a ready state for all servers: `racadm getremoteservicesstatus`
- PXE boot the target server.



### Why does the task: ‘Orchestrator: Deploy MetalLB IP Address pool’ fail?

```
TASK [orchestrator : Deploy MetalLB IP Address Pool] *****
task path: /root/test/devel-omnia/telemetry/roles/orchestrator/tasks/configure_metallb.yml:46
fatal: [localhost]: FAILED! => {"changed": true, "cmd": ["kubectl", "apply", "-f", "/var/lib/ipaddresspool.yaml"], "delta": "0:00:01.263525", "end": "2023-09-10 00:55:38.006362", "msg": "non-zero return code", "rc": 1, "start": "2023-09-10 00:55:36.742837", "stderr": "Error from server (InternalError): error when creating \"/var/lib/ipaddresspool.yaml\": Internal error occurred: failed calling webhook \"ipaddresspoolvalidationwebhook.metallb.io\": Post \"https://webhook-service.metallb-system.svc:443/validate-metallb-io-v1beta1-ipaddresspool?timeout=10s\": dial tcp 10.98.250.10:443: connect: connection refused\", \"stderr_lines\": [\"Error from server (InternalError): error when creating \"/var/lib/ipaddresspool.yaml\": Internal error occurred: failed calling webhook \"ipaddresspoolvalidationwebhook.metallb.io\": Post \"https://webhook-service.metallb-system.svc:443/validate-metallb-io-v1beta1-ipaddresspool?timeout=10s\": dial tcp 10.98.250.10:443: connect: connection refused\"], \"stdout\": \"\", \"stdout_lines\": []}
```

**Potential Cause:** /var partition is full (potentially due to images not being cleared after intel-oneapi images docker images are used to execute benchmarks on the cluster using aptainer support) .

**Resolution:** Clear the /var partition and retry telemetry.yml.

**Why does the task: [Telemetry]: TASK [grafana : Wait for grafana pod to come to ready state] fail with a timeout error?**

**Potential Cause:** Docker pull limit exceeded.

**Resolution:** Manually input the username and password to your docker account on the control plane.

### Is Disabling 2FA supported by Omnia?

- Disabling 2FA is not supported by Omnia and must be manually disabled.

### Is provisioning servers using BOSS controller supported by Omnia?

Provisioning server using BOSS controller is now supported by Omnia 1.2.1.

### How many IPs are required within the PXE NIC range?

Ensure that the number of IPs available between `pxe_nic_start_range` and `pxe_nic_end_range` is double the number of iDRACs available to account for potential stale entries in the mapping DB.

### What are the licenses required when deploying a cluster through Omnia?

While Omnia playbooks are licensed by Apache 2.0, Omnia deploys multiple softwares that are licensed separately by their respective developer communities. For a comprehensive list of software and their licenses, [click here](#) .

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

## 5.3 Troubleshooting guide

### 5.3.1 Connecting to internal databases

#### • TimescaleDB

- Go inside the pod: `kubectl exec -it pod/timescaledb-0 -n telemetry-and-visualizations -- /bin/bash`
- Connect to psql: `psql -U <postgres_username>`
- Connect to database: `\c telemetry_metrics`

#### • MySQL DB

- Go inside the pod: `kubectl exec -it pod/mysqlldb-0 -n telemetry-and-visualizations -- /bin/bash`

- Connect to mysql: `mysql -U <mysql_username> -p <mysql_password>`
- Connect to database: `USE idrac_telemetrysource_services_db`

### 5.3.2 Checking and updating encrypted parameters

1. Move to the filepath where the parameters are saved (as an example, we will be using `provision_config.yml`):

```
cd input/
```

2. To view the encrypted parameters:

```
ansible-vault view provision_config.yml --vault-password-file .provision_vault_key
```

3. To edit the encrypted parameters:

```
ansible-vault edit provision_config.yml --vault-password-file .provision_vault_key
```

### 5.3.3 Checking pod status on the control plane

- Select the pod you need to troubleshoot from the output of `kubectl get pods -A`
- Check the status of the pod by running `kubectl describe pod <pod name> -n <namespace name>`





### 5.3.4 Using telemetry information to diagnose node issues

Table 1: Regular telemetry metrics

| Metric Name             | Unit      | Possible Values   | Possible error causes   |
|-------------------------|-----------|---|---|
| BlockedProcesses        | processes | <ul style="list-style-type: none"> <li>Metric Value</li> <li>No Data</li> </ul> | <ul style="list-style-type: none"> <li>This could happen if the <code>/proc/stat</code> file is inaccessible.</li> </ul>    |
| CPUSystem               | seconds   | <ul style="list-style-type: none"> <li>Metric Value</li> <li>No Data</li> </ul> | <ul style="list-style-type: none"> <li>This could happen when the <code>psutil</code> library encounters errors.</li> </ul> |
| CPUWait                 | seconds   | <ul style="list-style-type: none"> <li>Metric Value</li> <li>No Data</li> </ul> | <ul style="list-style-type: none"> <li>This could happen when the <code>psutil</code> library encounters errors.</li> </ul> |
| ErrorsRecv              |           | <ul style="list-style-type: none"> <li>Metric Value</li> <li>No Data</li> </ul> | <ul style="list-style-type: none"> <li>This could happen when the <code>psutil</code> library encounters errors.</li> </ul> |
| ErrorsSent              |           | <ul style="list-style-type: none"> <li>Metric Value</li> <li>No Data</li> </ul> | <ul style="list-style-type: none"> <li>This could happen when the <code>psutil</code> library encounters errors.</li> </ul> |
| FailedJobs              |           | <ul style="list-style-type: none"> <li>Metric Value</li> <li>No Data</li> </ul> | <ul style="list-style-type: none"> <li>Slurm is not installed.</li> </ul>   |
| HardwareCorruptedMemory | kB        | <ul style="list-style-type: none"> <li>Metric Value</li> <li>No Data</li> </ul> | <ul style="list-style-type: none"> <li>This could happen if the <code>/proc/meminfo</code> file is inaccessible.</li> </ul> |
| MemoryActive            | bytes     | <ul style="list-style-type: none"> <li>Metric Value</li> <li>No Data</li> </ul> | <ul style="list-style-type: none"> <li>This could happen when the <code>psutil</code> library encounters errors.</li> </ul> |
| MemoryAvailable         | bytes     | <ul style="list-style-type: none"> <li>Metric Value</li> <li>No Data</li> </ul> | <ul style="list-style-type: none"> <li>This could happen when the <code>psutil</code> library encounters errors.</li> </ul> |
| MemoryCached            | bytes     | <ul style="list-style-type: none"> <li>Metric Value</li> <li>No Data</li> </ul> | <ul style="list-style-type: none"> <li>This could happen when the <code>psutil</code> library encounters errors.</li> </ul> |

### 5.3. Troubleshooting guide

|            |       |   |   |
|------------|-------|---|---|
| MemoryFree | bytes | <ul style="list-style-type: none"> <li>Metric Value</li> <li>No Data</li> </ul> | <ul style="list-style-type: none"> <li>This could happen when the <code>psutil</code> library encounters errors.</li> </ul> |
|------------|-------|---|---|

Table 2: Health telemetry metrics

| Metric Name                        | Possible value(s)   | Possible failure causes  |
|------------------------------------|---|--|
| dmesg                              | <ul style="list-style-type: none"> <li>• Unknown</li> <li>• Fail</li> <li>• Pass</li> </ul> | <ul style="list-style-type: none"> <li>• [Unknown] The dmesg command was not found on the cluster node.</li> <li>• [Fail] The dmesg command returned an error log message.</li> </ul>                                      |
| beegfs -beegfsstat                 | <ul style="list-style-type: none"> <li>• Unknown</li> <li>• Fail</li> <li>• Pass</li> </ul> | <ul style="list-style-type: none"> <li>• [Unknown] BeeGFS is not installed or inactive.</li> <li>• [Fail] The BeeGFS client service has failed or the node is not present in reachable lists of BeeGFS clients.</li> </ul> |
| gpu_driver_health:gpu              | <ul style="list-style-type: none"> <li>• Unknown</li> <li>• Fail</li> <li>• Pass</li> </ul> | <ul style="list-style-type: none"> <li>• AMD/NVIDIA accelerators are not present.</li> <li>• GPU drivers are not installed including Rocm and CUDA.</li> </ul>   |
| gpu_health_nvlink:gpu <sup>1</sup> | <ul style="list-style-type: none"> <li>• Unknown</li> <li>• Fail</li> <li>• Pass</li> </ul> | <ul style="list-style-type: none"> <li>• AMD/NVIDIA accelerators are not present.</li> <li>• NVLinks are not present.</li> <li>• GPU drivers are not installed including Rocm and CUDA.</li> </ul>                         |
| gpu_health_pcie:gpu                | <ul style="list-style-type: none"> <li>• Unknown</li> <li>• Fail</li> <li>• Pass</li> </ul> | <ul style="list-style-type: none"> <li>• AMD/NVIDIA accelerators are not present.</li> <li>• GPU drivers are not installed including Rocm and CUDA.</li> </ul>   |
| gpu_health_pmu:gpu                 | <ul style="list-style-type: none"> <li>• Unknown</li> <li>• Fail</li> <li>• Pass</li> </ul> | <ul style="list-style-type: none"> <li>• AMD/NVIDIA accelerators are not present.</li> <li>• GPU drivers are not installed including Rocm and CUDA.</li> </ul>   |
| gpu_health_power:gpu               | <ul style="list-style-type: none"> <li>• Unknown</li> <li>• Fail</li> <li>• Pass</li> </ul> | <ul style="list-style-type: none"> <li>• AMD/NVIDIA accelerators are not present.</li> <li>• GPU drivers are not installed including Rocm and CUDA.</li> </ul>   |
| gpu_health_thermal:gpu             | <ul style="list-style-type: none"> <li>• Unknown</li> <li>• Metric Value</li> </ul>         | <ul style="list-style-type: none"> <li>• AMD/NVIDIA accelerators are not present.</li> <li>• GPU drivers are not installed including Rocm and CUDA.</li> </ul>   |
| Kubernetespodsstatus               | <ul style="list-style-type: none"> <li>• Unknown</li> <li>• Fail</li> <li>• Pass</li> </ul> | <ul style="list-style-type: none"> <li>• Kubernetes is not installed.</li> </ul>   |
| <b>178</b>                         |   | <b>Chapter 5. Troubleshooting</b>  |
| Kuberneteschildnode                | <ul style="list-style-type: none"> <li>• Unknown</li> <li>• Fail</li> <li>• Pass</li> </ul> | <ul style="list-style-type: none"> <li>• Kubernetes is not installed.</li> </ul>   |

Table 3: GPU telemetry metrics

| Metric Name             | Unit    | Possible value(s)   | Potential error cause(s)   |
|-------------------------|---------|---|--|
| gpu_temperature:gpu     | C       | <ul style="list-style-type: none"> <li>• Metric value</li> <li>• No data</li> </ul> | <ul style="list-style-type: none"> <li>• AMD/NVIDIA accelerators are not present.</li> <li>• GPU drivers are not installed including Rocm and CUDA.</li> </ul> |
| gpu_utilization         | percent | <ul style="list-style-type: none"> <li>• Metric value</li> <li>• No data</li> </ul> | <ul style="list-style-type: none"> <li>• AMD/NVIDIA accelerators are not present.</li> <li>• GPU drivers are not installed including Rocm and CUDA.</li> </ul> |
| gpu_utilization:average | percent | <ul style="list-style-type: none"> <li>• Metric value</li> <li>• No data</li> </ul> | <ul style="list-style-type: none"> <li>• AMD/NVIDIA accelerators are not present.</li> <li>• GPU drivers are not installed including Rocm and CUDA.</li> </ul> |

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

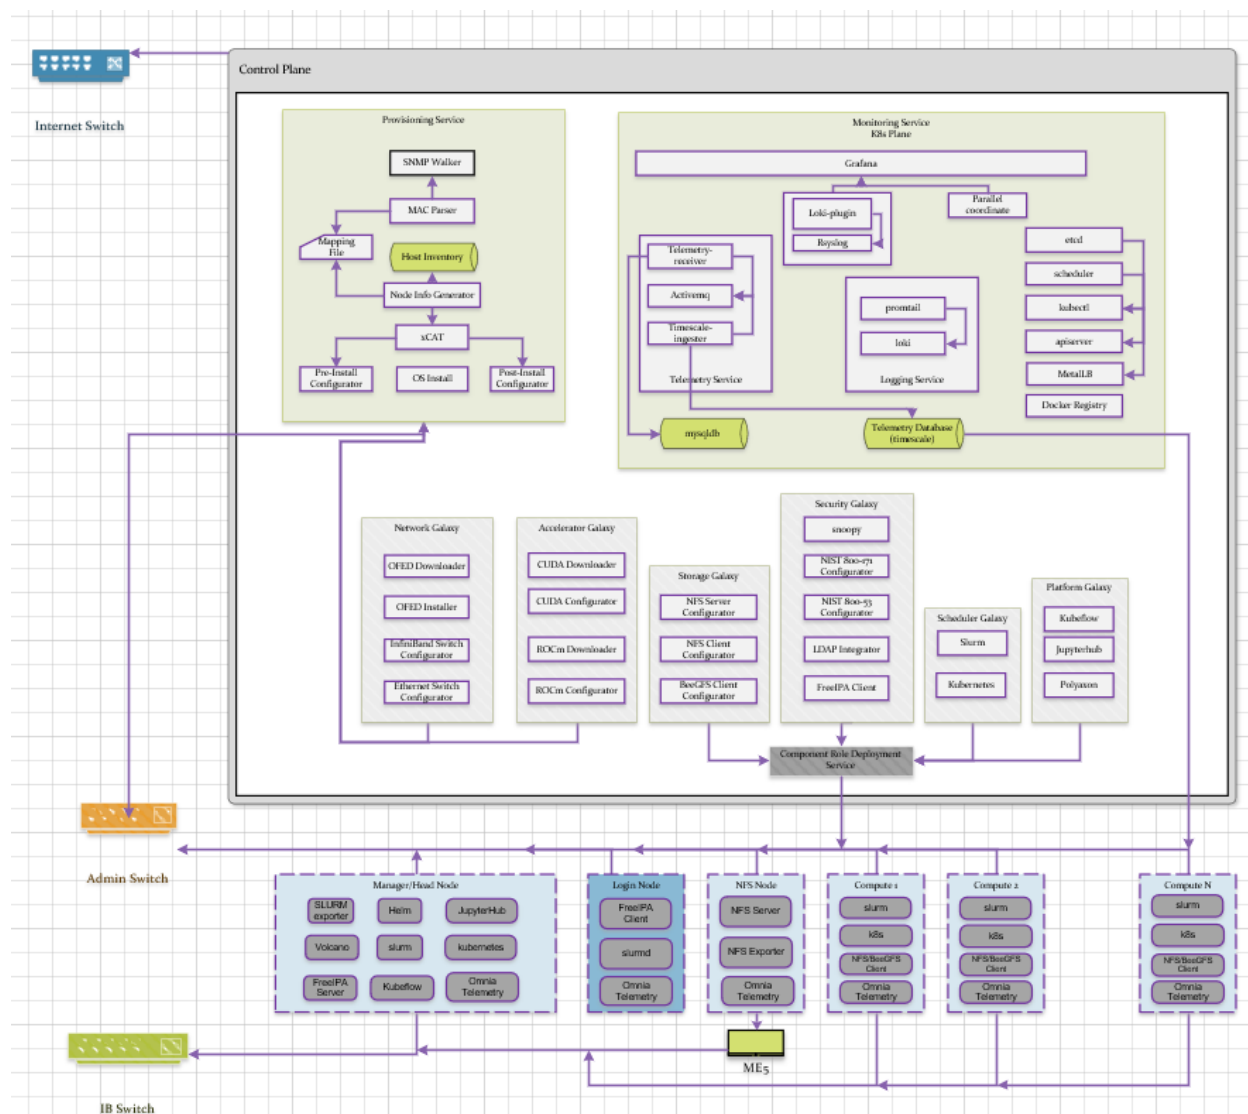
<sup>1</sup> This metric is collected from the manager node if a login node is absent.



## **SECURITY CONFIGURATION GUIDE**

### **6.1 Preface**

The security configuration guide of Omnia provides Dell customers an overview and understanding of the security features supported by Omnia 1.4. As part of an effort to improve its product lines, Dell periodically releases revisions of its software and hardware. The product release notes provide the most up-to-date information about product features. Contact your Dell technical support professional if a product does not function properly or does not function as described in this document. This document was accurate at publication time. To ensure that you are using the latest version of this document, go to [Omnia: Docs](#).



### 6.1.1 Legal disclaimers

THE INFORMATION IN THIS PUBLICATION IS PROVIDED “AS-IS.” DELL MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. In no event shall Dell Technologies, its affiliates or suppliers, be liable for any damages whatsoever arising from or related to the information contained herein or actions that you decide to take based thereon, including any direct, indirect, incidental, consequential, loss of business profits or special damages, even if Dell Technologies, its affiliates or suppliers have been advised of the possibility of such damages. The Security Configuration Guide intends to be a reference. The guidance is provided based on a diverse set of installed systems and may not represent the actual risk/guidance to your local installation and individual environment. It is recommended that all users determine the applicability of this information to their individual environments and take appropriate actions. All aspects of this Security Configuration Guide are subject to change without notice and on a case-by-case basis. Your use of the information contained in this document or materials linked herein is at your own risk. Dell reserves the right to change or update this document in its sole discretion and without notice at any time.

## 6.1.2 Scope of the document

This document covers the security features supported by Omnia 1.4.

## 6.1.3 Document references

In addition to this guide, more information on Omnia can be found through the below links:

- [Omnia: Read Me](#)
- [Omnia: Quick Installation Guide](#)

## 6.1.4 Reporting security vulnerabilities

Dell takes reports of potential security vulnerabilities in our products very seriously. If you discover a security vulnerability, you are encouraged to report it to Dell immediately. For the latest instructions on how to report a security issue to Dell, see the [Dell Vulnerability Response Policy](#) on the Dell.com site.

Follow Dell Security on these sites:

- [dell.com/security](https://dell.com/security)
- [dell.com/support](https://dell.com/support)

To provide feedback on this solution, email us at [support@dell.com](mailto:support@dell.com).

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

# 6.2 Security Quick Reference


## 6.2.1 Security profiles

Omnia requires root privileges during installation because it provisions the operating system on bare metal servers.

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

## 6.3 Product and Subsystem Security

### 6.3.1 Security controls map



Omnia performs bare metal configuration to enable AI/HPC workloads. It uses Ansible playbooks to perform installations and configurations. iDRAC is supported for provisioning bare metal servers. Omnia installs xCAT to enable provisioning of clusters via PXE in different ways:

- Mapping file **[optional]**: To dictate IP address/MAC mapping, a host mapping file can be provided.
- BMC discovery **[optional]**: To discover the cluster via BMC (iDRAC), IPMI must be enabled on remote servers. Discovery happens over IPMI. For security best practices when using this method, [click here!](#)
- SNMP **[optional]**: To discover the cluster by querying switches, SNMPv2 must be enabled.
- Switch **[default]**: To discovery the cluster by routing communication through particular switch ports over SNMPv3, non-admin switch credentials must be provided.

**Note:** IPMI is not required on the control plane. However compute nodes (iDRACs in the cluster/private network) require IPMI to be enabled for BMC discovery.

Omnia can be installed via CLI only. Slurm and Kubernetes are deployed and configured on the cluster. FreeIPA or LDAP is installed for providing authentication.



To perform these configurations and installations, a secure SSH channel is established between the management node and the following entities:

- Manager Node
- Compute Nodes
- Login Node

### 6.3.2 Authentication

Omnia does not have its own authentication mechanism because bare metal installations and configurations take place using root privileges. Post the execution of Omnia, third-party tools are responsible for authentication to the respective tool.

### 6.3.3 Cluster authentication tool

In order to enable authentication to the cluster, Omnia installs FreeIPA: an open source tool providing integrated identity and authentication for Linux/UNIX networked environments. As part of the HPC cluster, the login node is responsible for configuring users and managing a limited number of administrative tasks. Access to the manager/head node is restricted to administrators with the root password. For authentication on the manager and compute nodes exclusively, LDAP can also be installed by Omnia on the client.

---

**Note:** Omnia does not configure LDAP users or groups.

---

### 6.3.4 Authentication types and setup

#### Key-Based authentication

##### Use of SSH authorized\_keys

A password-less channel is created between the management station and compute nodes using SSH authorized keys. This is explained in Security Controls Maps.

### 6.3.5 Login security settings

The following credentials have to be entered to enable different tools on the management station:

1. iDRAC (Username/ Password)
2. Ethernet Switch (Username/ Password)
3. Infiniband Switch (Username/ Password)
4. PowerVault ME4/ME5 (Username/ Password)
5. Provisioning OS (Password)
6. SNMPv3 PXE switch (Non-admin username/ password)

Similarly, passwords for the following tools have to be provided in `input/omnia_config.yml` to configure the cluster:

1. maria\_db (Password)
2. DockerHub (Username/ Password)

For setting up authentication on the cluster, the following credentials have to be provided in `input/security_config.yml`:

1. FreeIPA (`directory_manager_password`, `ipa_admin_password`)
2. LDAP (`ldap_bind_username`, `ldap_bind_password`)

Once Omnia is invoked, these files are validated and encrypted using Ansible Vault. They are hidden from external visibility and access.

## 6.4 Authentication to external systems

Third party software installed by Omnia are responsible for supporting and maintaining manufactured-unique or installation-unique secrets.

### 6.4.1 Configuring remote connections

When setting up BeeGFS client services on the cluster, a connection authentication file is used to maintain the security of the communications between server and client.

1. Generate the connection authentication file (`connAuth`) and use it to set up BeeGFS meta, server and storage services.
2. Copy the `connAuth` file to the control plane and note the filepath.
3. Populate the value of `beegfs_secret_storage_filepath` in `input/storage_config.yml` with the filepath from the previous step.

Omnia will configure the BeeGFS clients on the cluster using the provided file. BeeGFS is responsible for maintaining and securing `connAuthFile`. For more information, [click here](#).

## 6.5 Network security

Omnia configures the firewall as required by the third-party tools to enhance security by restricting inbound and outbound traffic to the TCP and UDP ports.

### 6.5.1 Network exposure

Omnia uses port 22 for SSH connections as Ansible uses port 22.

### 6.5.2 Firewall settings

Omnia configures the following ports for use by third-party tools installed by Omnia.

#### **Kubernetes ports requirements**

| Port        | Number | Layer 4                 | Protocol Purpose | Type of Node |
|-------------|--------|-------------------------|------------------|--------------|
| 6443        | TCP    | Kubernetes API          | server           | Manager      |
| 2379-2380   | TCP    | etcd server             | client           | API Manager  |
| 10251       | TCP    | Kube-scheduler          | Manager          |              |
| 10252       | TCP    | Kube-controller manager | Manager          |              |
| 10250       | TCP    | Kubelet API             | Compute          |              |
| 30000-32767 | TCP    | Nodeport services       | Compute          |              |
| 5473        | TCP    | Calico services         | Manager/Compute  |              |
| 179         | TCP    | Calico services         | Manager/Compute  |              |
| 4789        | UDP    | Calico services         | Manager/Compute  |              |
| 8285        | UDP    | Flannel services        | Manager/Compute  |              |
| 8472        | UDP    | Flannel services        | Manager/Compute  |              |

### Slurm port requirements

| Port | Number  | Layer 4        | Protocol | Node |
|------|---------|----------------|----------|------|
| 6817 | TCP/UDP | Slurmctld Port | Manager  |      |
| 6818 | TCP/UDP | Slurmd Port    | Compute  |      |
| 6819 | TCP/UDP | Slurmdbd Port  | Manager  |      |

### BeeGFS port requirements

| Port | Service                          |
|------|----------------------------------|
| 8008 | Management service (beegfs-mgmt) |
| 8003 | Storage service (beegfs-storage) |
| 8004 | Client service (beegfs-client)   |
| 8005 | Metadata service (beegfs-meta)   |
| 8006 | Helper service (beegfs-helper)   |

### xCAT port requirements

| Port number   | Protocol | Service Name |
|---------------|----------|--------------|
| 3001          | tcp      | xcatdport    |
| 3001          | udp      | xcatdport    |
| 3002          | tcp      | xcatiport    |
| 3002          | udp      | xcatiport    |
| 3003(default) | tcp      | xcatlport    |
| 7             | udp      | echo-udp     |
| 22            | tcp      | ssh-tcp      |
| 22            | udp      | ssh-udp      |
| 873           | tcp      | rsync        |
| 873           | udp      | rsync        |
| 53            | tcp      | domain-tcp   |
| 53            | udp      | domain-udp   |
| 67            | udp      | bootps       |
| 67            | tcp      | dhcp         |
| 68            | tcp      | dhcpc        |

continues on next page

Table 1 – continued from previous page

| Port number | Protocol | Service Name |
|-------------|----------|--------------|
| 68          | udp      | bootpc       |
| 69          | tcp      | tftp-tcp     |
| 69          | udp      | tftp-udp     |
| 80          | tcp      | www-tcp      |
| 80          | udp      | www-udp      |
| 88          | tcp      | kerberos     |
| 88          | udp      | kerberos     |
| 111         | udp      | sunrpc-udp   |
| 443         | udp      | HTTPS        |
| 443         | tcp      | HTTPS        |
| 514         | tcp      | shell        |
| 514         | tcp      | rsyslogd     |
| 514         | udp      | rsyslogd     |
| 544         | tcp      | kshell       |
| 657         | tcp      | rmc-tcp      |
| 657         | udp      | rmc-udp      |
| 782         | tcp      | conserver    |
| 1058        | tcp      | nim          |
| 2049        | tcp      | nfsd-tcp     |
| 2049        | udp      | nfsd-udp     |
| 4011        | tcp      | pxe          |
| 300         | tcp      | awk          |
| 623         | tcp      | ipmi         |
| 623         | udp      | ipmi         |
| 161         | tcp      | snmp         |
| 161         | udp      | snmp         |
| 162         | tcp      | snmptrap     |
| 162         | udp      | snmptrap     |
| 5432        | tcp      | postgresDB   |

**Note:** For more information, check out the [xCAT website](#).

### FreeIPA port requirements

| Port Number | Layer 4 | Purpose              | Node                |
|-------------|---------|----------------------|---------------------|
| 80          | TCP     | HTTP/HTTPS           | Manager/ Login_Node |
| 443         | TCP     | HTTP/HTTPS           | Manager/ Login_Node |
| 389         | TCP     | LDAP/LDAPS           | Manager/ Login_Node |
| 636         | TCP     | LDAP/LDAPS           | Manager/ Login_Node |
| 88          | TCP/UDP | Kerberos             | Manager/ Login_Node |
| 464         | TCP/UDP | Kerberos             | Manager/ Login_Node |
| 53          | TCP/UDP | DNS                  | Manager/ Login_Node |
| 7389        | TCP     | Dogtag's LDAP server | Manager/ Login_Node |
| 123         | UDP     | NTP                  | Manager/ Login_Node |

**Note:** To avoid security vulnerabilities, protocols can be restricted on the network using the parameters re-

`strict_program_support` and `restrict_softwares`. However, certain protocols are essential to Omnia's functioning and cannot be disabled: `ftp`, `smbd`, `nmbd`, `automount`, `portmap`.

### 6.5.3 Data security

Omnia does not store data. The passwords Omnia accepts as input to configure the third party tools are validated and then encrypted using Ansible Vault. Run `yum update --security` routinely on the control plane for the latest security updates.

For more information on the passwords used by Omnia, see Login Security Settings.

### 6.5.4 Auditing and logging

Omnia creates a log file at `/var/log/omnia` on the management station. The events during the installation of Omnia are captured as logs. For different roles called by Omnia, separate log files are created as listed below:

- `monitor.log`
- `network.log`
- `provision.log`
- `scheduler.log`
- `security.log`
- `storage.log`
- `utils.log`

Additionally, an aggregate of the events taking place during storage, scheduler and network role installation called `omnia.log` is created in `/var/log`.

There are separate logs generated by the third party tools installed by Omnia.

### 6.5.5 Logs

A sample of the `omnia.log` is provided below:

```
2021-02-15 15:17:36,877 p=2778 u=omnia n=ansible | [WARNING]: provided hosts
list is empty, only localhost is available. Note that the implicit localhost does not
match 'all'
2021-02-15 15:17:37,396 p=2778 u=omnia n=ansible | PLAY [Executing omnia roles]
*****
2021-02-15 15:17:37,454 p=2778 u=omnia n=ansible | TASK [Gathering Facts]
*****
*
2021-02-15 15:17:38,856 p=2778 u=omnia n=ansible | ok: [localhost]
2021-02-15 15:17:38,885 p=2778 u=omnia n=ansible | TASK [common : Mount Path]
*****
2021-02-15 15:17:38,969 p=2778 u=omnia n=ansible | ok: [localhost]
```

These logs are intended to enable debugging.

**Note:** The Omnia product recommends that product users apply masking rules on personal identifiable information (PII) in the logs before sending to external monitoring applications or sources.

## 6.5.6 Logging format

Every log message begins with a timestamp and also carries information on the invoking play and task.

The format is described in the following table.

| Field                            | Format                           | Sample Value  |
|----------------------------------|----------------------------------|---|
| Timestamp                        | yyyy-mm-dd h:m:s                 | 2/15/2021 15:17   |
| Process Id                       | p=xxxx                           | p=2778  |
| User                             | u=xxxx                           | u=omnia   |
| Name of the process executing    | n=xxxx                           | n=ansible   |
| The task being executed/ invoked | PLAY/TASK                        | PLAY [Executing omnia roles] TASK [Gathering Facts]           |
| Error                            | fatal: [hostname]: Error Message | fatal: [localhost]: FAILED! => {"msg": "lookup_plugin.lines"} |
| Warning                          | [WARNING]: warning message       | [WARNING]: provided hosts list is empty                       |

## 6.5.7 Network vulnerability scanning

Omnia performs network security scans on all modules of the product. Omnia additionally performs Blackduck scans on the open source softwares, which are installed by Omnia at runtime. However, Omnia is not responsible for the third-party software installed using Omnia. Review all third party software before using Omnia to install it.

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

## 6.6 Miscellaneous Configuration and Management Elements

### 6.6.1 Licensing

Omnia 1.4 is licensed under the Apache License 2.0. A permissive license whose main conditions require preservation of copyright and license notices. Contributors provide an express grant of patent rights. Licensed works, modifications, and larger works may be distributed under different terms and without source code.

## 6.6.2 Protect authenticity

Every GitHub push requires a sign-off and a moderator is required to approve pull requests. All contributions have to be certified using the Developer Certificate of Origin (DCO):

Developer Certificate of Origin  
Version 1.1

Copyright (C) 2004, 2006 The Linux Foundation and its contributors.  
1 Letterman Drive  
Suite D4700  
San Francisco, CA, 94129

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

Developer's Certificate of Origin 1.1

By making a contribution to this project, I certify that:

- (a) The contribution was created in whole or in part by me and I have the right to submit it under the open source license indicated in the file; or
- (b) The contribution is based upon previous work that, to the best of my knowledge, is covered under an appropriate open source license and I have the right under that license to submit that work with modifications, whether created in whole or in part by me, under the same open source license (unless I am permitted to submit under a different license), as indicated in the file; or
- (c) The contribution was provided directly to me by some other person who certified (a), (b) or (c) and I have not modified it.
- (d) I understand and agree that this project and the contribution are public and that a record of the contribution (including all personal information I submit with it, including my sign-off) is maintained indefinitely and may be redistributed consistent with this project or the open source license(s) involved.

### 6.6.3 Ansible security

For the security guidelines of Ansible modules, go to [Developing Modules Best Practices: Module Security](#).

### 6.6.4 Ansible vault

Ansible vault enables encryption of variables and files to protect sensitive content such as passwords or keys rather than leaving it visible as plaintext in playbooks or roles. Please refer [Ansible Vault guidelines](#) for more information.

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).



## SAMPLE FILES

**Caution:** All the file contents mentioned below are case sensitive. The casing of words like [manager], [compute], etc should be consistent with the samples below when creating inventory or mapping files.

### 7.1 inventory file

```
[manager]
10.5.0.101

[compute]
10.5.0.102
10.5.0.103

[login]
10.5.0.104
```

### 7.2 pxe\_mapping\_file.csv

```
MAC,Hostname,IP
xx:yy:zz:aa:bb:cc,server,10.5.0.101
aa:bb:cc:dd:ee:ff,server2, 10.5.0.102
```

---

**Note:**

- To skip the provisioning of a particular node in the list, simply append a '#' to the beginning of the line pertaining to that node.
  - Hostnames listed in this file should be exclusively lower-case with no special characters.
-

## 7.3 switch\_inventory

```
10.3.0.101  
10.3.0.102
```

## 7.4 powervault\_inventory

```
10.3.0.105
```

## 7.5 NFS Server inventory file

```
[nfs]  
10.5.0.104
```

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

## **LIMITATIONS**

- Once `provision.yml` is used to configure devices, it is recommended to avoid rebooting the control plane.
- Omnia provision tools only support /16 subnet masks for provisioning.
- Omnia supports adding only 1000 nodes when discovered via BMC.
- Removal of Slurm and Kubernetes component roles are not supported. However, the scheduler type can be customized by setting `scheduler_type` in `input/omnia_config.yml` prior to running `omnia.yml`.
- After installing the Omnia control plane, changing the manager node is not supported. If you need to change the manager node, you must redeploy the entire cluster.
- Dell Technologies provides support to the Dell-developed modules of Omnia. All the other third-party tools deployed by Omnia are outside the support scope.
- To change the Kubernetes single node cluster to a multi-node cluster or change a multi-node cluster to a single node cluster, you must either redeploy the entire cluster or run `kubeadm reset -f` on all the nodes of the cluster. Then set `scheduler_type:k8s` in `input/omnia_config.yml` prior to running `omnia.yml`.
- In a single node cluster, the login node and Slurm functionalities are not applicable. However, Omnia installs FreeIPA Server and Slurm on the single node.
- To change the Kubernetes version from 1.16 to 1.19 or 1.19 to 1.16, you must redeploy the entire cluster.
- The Kubernetes pods will not be able to access the Internet or start when `firewalld` is enabled on the node. This is a limitation in Kubernetes. So, the `firewalld` daemon will be disabled on all the nodes as part of `omnia.yml` execution.
- Only one storage instance (Powervault) is currently supported in the HPC cluster.
- Omnia supports only basic telemetry configurations. Changing data fetching time intervals for telemetry is not supported.
- Slurm cluster metrics will only be fetched from clusters configured by Omnia.
- All iDRACs must have the same username and password.
- OpenSUSE Leap 15.3 is not supported on the Control Plane.
- Omnia might contain some unused MACs since LOM switch have both iDRAC MACs as well as ethernet MACs, PXE NIC ranges should contain IPs that are double the iDRACs present.
- FreeIPA authentication is not supported on the control plane.
- The multiple OS feature is only available with Rocky 8.7 when xCAT 2.16.5 is in use. Currently, Omnia uses 2.16.4.
- Currently, Omnia only supports the splitting of switch ports. Switch ports cannot be un-split using this script.

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).



## **BEST PRACTICES**

- Ensure that PowerCap policy is disabled and the BIOS system profile is set to 'Performance' on the Control Plane.
- Ensure that there is at least 50% (~50GB) free space on the Control Plane root partition before running Omnia. To maintain the free space required, place any ISO files used in the /home directory.
- Disable SELinux on the Control Plane.
- Use a [PXE mapping file](#) even when using DHCP configuration to ensure that IP assignments remain persistent across Control Plane reboots.
- Avoid rebooting the Control Plane as much as possible to ensure that all network configuration does not get disturbed.
- Review the prerequisites before running Omnia Scripts.
- Ensure that the firefox version being used on the control plane is the latest available. This can be achieved using `dnf update firefox -y`
- It is recommended to configure devices using Omnia playbooks for better interoperability and ease of access.
- Ensure that the /var partition has adequate space to complete commands.
- Run `yum update --security` routinely on the control plane for the latest security updates.

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).



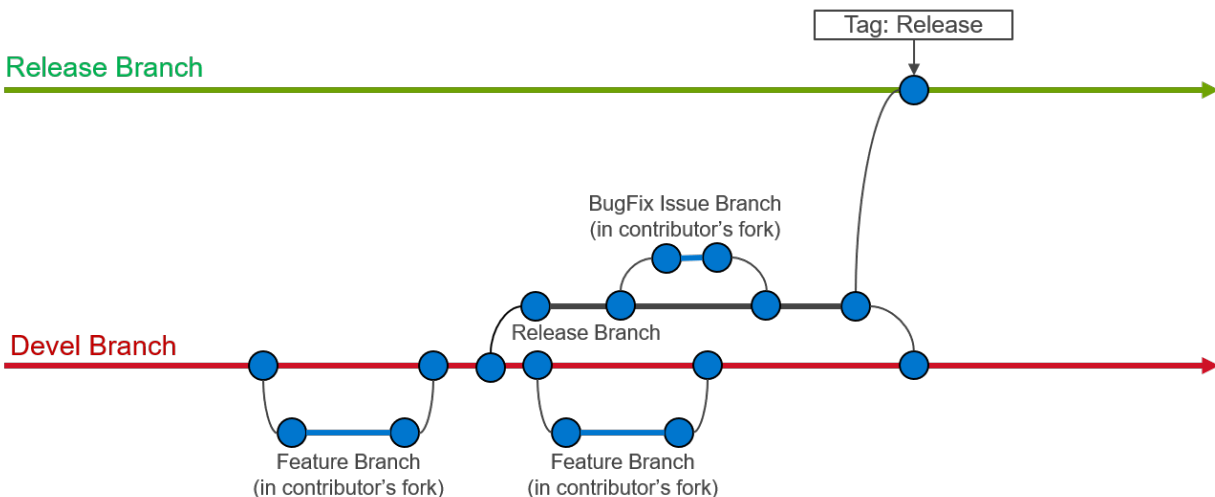
## CONTRIBUTING TO OMNIA

We encourage everyone to help us improve Omnia by contributing to the project. Contributions can be as small as documentation updates or adding example use cases, to adding commenting and properly styling code segments all the way up to full feature contributions. We ask that contributors follow our established guidelines for contributing to the project.

This document will evolve as the project matures. Please be sure to regularly refer back in order to stay in-line with contribution guidelines.

### 10.1 Creating A Pull Request

Contributions to Omnia are made through [Pull Requests \(PRs\)](#). To make a pull request against Omnia, use the following steps.



### 10.1.1 Create an issue

Create an [issue](#) and describe what you are trying to solve. It does not matter whether it is a new feature, a bug fix, or an improvement. All pull requests must be associated to an issue. When creating an issue, be sure to use the appropriate issue template (bug fix or feature request) and complete all of the required fields. If your issue does not fit in either a bug fix or feature request, then create a blank issue and be sure to including the following information:

- **Problem description:** Describe what you believe needs to be addressed
- **Problem location:** In which file and at what line does this issue occur?
- **Suggested resolution:** How do you intend to resolve the problem?

### 10.1.2 Fork the repository

All work on Omnia should be done in a [fork of the repository](#). Only maintainers are allowed to commit directly to the project repository.

### 10.1.3 Issue branch

Create a [new branch](#) on your fork of the repository. All contributions should be branched from `devel`:

```
git checkout devel
git checkout -b <new-branch-name>
```

**Branch name:** The branch name should be based on the issue you are addressing. Use the following pattern to create your new branch name: `issue-xxxx`, e.g., `issue-1023`.

### 10.1.4 Commit changes

- It is important to commit your changes to the issue branch. Commit messages should be descriptive of the changes being made.
- All commits to Omnia need to be signed with the [Developer Certificate of Origin \(DCO\)](#) in order to certify that the contributor has permission to contribute the code. In order to sign commits, use either the `--signoff` or `-s` option to `git commit`:

```
git commit --signoff
git commit -s
```

Make sure you have your user name and e-mail set. The `--signoff` | `-s` option will use the configured user name and e-mail, so it is important to configure it before the first time you commit. Check the following references:

- [Setting up your github user name](#)
- [Setting up your e-mail address](#)

**Caution:** When preparing a pull request it is important to stay up-to-date with the project repository. We recommend that you rebase against the upstream repo frequently.

```
git pull --rebase upstream devel #upstream is dell/omnia
git push --force origin <pr-branch-name> #origin is your fork of the repository (e.g.,
↪ <github_user_name>/omnia.git)
```



### 10.1.5 PR description

Be sure to fully describe the pull request. Ideally, your PR description will contain:

1. A description of the main point (i.e., why was this PR made?),
2. Linking text to the related issue (i.e., This PR closes issue #<issue\_number>),
3. How the changes solves the problem
4. How to verify that the changes work correctly.

### 10.1.6 Developer Certificate of Origin

```
Developer Certificate of Origin
Version 1.1

Copyright (C) 2004, 2006 The Linux Foundation and its contributors.
1 Letterman Drive
Suite D4700
San Francisco, CA, 94129

Everyone is permitted to copy and distribute verbatim copies of this
license document, but changing it is not allowed.

Developer's Certificate of Origin 1.1

By making a contribution to this project, I certify that:

(a) The contribution was created in whole or in part by me and I
    have the right to submit it under the open source license
    indicated in the file; or

(b) The contribution is based upon previous work that, to the best
    of my knowledge, is covered under an appropriate open source
    license and I have the right under that license to submit that
    work with modifications, whether created in whole or in part
    by me, under the same open source license (unless I am
    permitted to submit under a different license), as indicated
    in the file; or

(c) The contribution was provided directly to me by some other
    person who certified (a), (b) or (c) and I have not modified
    it.

(d) I understand and agree that this project and the contribution
    are public and that a record of the contribution (including all
    personal information I submit with it, including my sign-off) is
    maintained indefinitely and may be redistributed consistent with
    this project or the open source license(s) involved.
```

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).

If you have any feedback about Omnia documentation, please reach out at [omnia.readme@dell.com](mailto:omnia.readme@dell.com).