

# Powering Up a New SGI® Octane™ III Deskside Cluster Server

---

Powering up a new SGI Octane III deskside cluster server requires powering up the headnode and the compute nodes. The process to do this depends on the configuration that was purchased. This document describes the default configurations of newly shipped SGI Octane III deskside clusters and provides procedures for powering up a new system.

## Possible Configurations

SGI Octane III deskside clusters ship from SGI manufacturing with one of the configurations shown in Table 1.

**Table 1** Possible Configurations

Configuration	Description
No OS installed and no factory integration performed	<p>The cluster hardware has been configured and tested. No operating system or cluster manager is installed on the headnode or compute nodes.</p> <p>You must install an operating system, cluster management software, and compute node images to make the system an operational cluster.</p> <p>You cannot power up the compute nodes and remotely manage them until you install an operating system on the headnode. If the selected operating includes IPMI support, you can use it to perform compute node power up and remote management.</p>
Linux® OS installed but no factory integration performed	<p>The cluster hardware has been configured and tested. The Linux operating system has been installed on the headnode. Nothing is installed on the compute nodes.</p> <p>You must install your own cluster management software and load your own compute node images to make the system an operational cluster.</p> <p>Use IPMI to power up and perform remote management of the compute nodes (<code>ipmitool</code> is installed on the headnode).</p>
Linux OS installed and factory integration performed (SGI Management Center*)	<p>The cluster hardware has been configured and tested. The requested Linux operating system and SGI Management Center cluster management software has been installed and configured for all nodes in the purchased configuration.</p> <p>Use the SGI Management Center graphical user interface, the SGI Management Center command-line interface, or IPMI to power up and perform remote management of the compute nodes (<code>ipmitool</code> is installed on the nodes).</p>
Microsoft® Windows® HPC OS installed but no factory integration performed	<p>The cluster hardware has been configured and tested. The operating system is installed on the headnode. You must run the <code>HPCGettingStarted</code> executable to install and configure the compute nodes.</p>

---

\* The SGI Management Center was formerly called the SGI ISLE™ Cluster Manager.

## Default Network Configuration

SGI Octane III systems ship from SGI manufacturing with the network configuration shown in Table 2.

**Table 2** Default Network Configuration\*

Component	Hostname	IP Address
<b>Headnode</b>		
nic1 (public)	Factory hostname**	Factory IP address**
nic2 (management GigE)	n/a	10.0.10.1
InfiniBand HCA	n/a	192.168.20.1
nic3 (optional GigE)	n/a	172.16.10.1
BMC	n/a	Factory IP address**
<b>Compute nodes</b>		
nic1 (management GigE)	n/a	10.0.1.1 10.0.1.2 ... 10.0.1.X
InfiniBand HCA	n/a	192.168.22.1 192.168.22.2 ... 192.168.22.X
nic2 (optional GigE)	n/a	172.16.1.1 172.16.1.2 ... 172.16.1.X
BMC	n001-bmc n002-bmc ... nXXX-bmc	10.0.40.1 10.0.40.2 ... 10.0.40.X
<b>Graphic nodes</b>		
nic1 (management GigE)	n/a	10.0.100.1 10.0.100.2 ... 10.0.100.X
InfiniBand HCA	n/a	192.168.100.1 192.168.100.2 ... 192.168.100.X

**Table 2 (continued)** Default Network Configuration\*

Component	Hostname	IP Address
<b>Graphic nodes (cont.)</b>		
nic2 (optional GigE)	n/a	172.16.100.1 172.16.100.2 ... 172.16.100.X
BMC	n/a	10.0.50.1 10.0.50.2 ... 10.0.50.X
<b>GigE switch</b>	n/a	10.0.20.1
<b>InfiniBand switch</b>	n/a	10.0.21.1

\* Specific settings for a system are included in the *SGI Cluster Network Configuration/Validation* document that ships with the system.

\*\* You must change the factory hostnames and IP addresses to your desired configuration.

## Default Login and Password Settings

SGI Octane III desktide clusters ship from SGI manufacturing with the login and password settings shown in Table 2.

**Table 3** Default Login and Password Settings

Component	User	Password
Headnode and compute nodes	root	sgisgi
GigE switch*	admin	
InfiniBand switch	admin	123456
ipmitool	admin	admin
SGI Management Center GUI (ISLE)	root	root

\* There is no password set on the GigE switch.

## Powering Up a New System

The procedures for powering up a system vary depending on the operating system and cluster management software that are installed.

### Powering Up a System with No Operating System Installed

Use the following procedure to power up a system that does not have an operating system installed and was not factory integrated:

**Note:** You must install an operating system and cluster management software on a system that ships without an operating system installed before the system will be an operational cluster. Before attempting to boot the system, insert a DVD into a USB DVD drive attached to the headnode or configure the headnode for network installation to install your desired software on the system.

1. Verify that the power cords are connected to the rear of the system and the site power source.
2. Move the system enable switch on the rear of the enclosure to the ON position

**Note:** The system enable switch applies power to all internal trays and switches. The headnode (located in the top tray) is the only node that is configured to power-on automatically. The compute nodes do not power-on automatically; you must start them manually with a cluster manager or IPMI if the operating system you install supports it.

3. Install the desired operating system and cluster management software on the headnode. Then, install the desired images on the compute nodes.

### Powering Up a System with Only the Linux OS Installed

Use the following procedure to power up a system that has the Linux OS installed but was not factory integrated:

1. Verify that the power cords are connected to the rear of the system and the site power source.
2. Move the system enable switch on the rear of the enclosure to the ON position

**Note:** The system enable switch applies power to all internal trays and switches. The headnode (located in the top tray) is the only node that is configured to power-on automatically. The compute nodes do not power-on automatically; you must start them manually with the `ipmitool` command.

3. Power up the compute nodes using the `ipmitool` command from the headnode:

```
ipmitool -I lanplus -H [node_bmc_ipaddr] -U admin -P admin power on
```

For example, the following command powers up compute node 1:

```
ipmitool -I lanplus -H 10.0.40.1-U admin -P admin power on
```

**Note:** You must install cluster management software and load your own compute node images to make the system an operational cluster.

## Powering Up a System with the Linux OS and SGI Management Center Installed

Use the following procedure to power up a system that has the Linux OS installed and was factory integrated with the SGI Management Center cluster management software:

1. Verify that the power cords are connected to the rear of the system and the site power source.
2. Move the system enable switch on the rear of the enclosure to the ON position

**Note:** The system enable switch applies power to all internal trays and switches. The headnode (located in the top tray) is the only node that is configured to power-on automatically. The compute nodes do not power-on automatically; you must start them manually with the cluster manager or `ipmitool` command.

3. Power up the compute nodes using the cluster management software or the `ipmitool` command:
  - To use the SGI Management Console graphical user interface, refer to the *SGI Management Center System Administrator's Guide*, publication number 007-5642-00x.
  - To use the SGI Management Console command-line interface, enter the following command from the headnode:

```
powerman -l -n [n001-00X]
```

For example, the following command powers up nodes 001 through 005:

```
powerman -l -n [n001-005]
```
  - To use `ipmitool`, enter the following command for each compute node from the headnode:

```
ipmitool -I lanplus -H [node_bmc_ipaddr] -U admin -P admin power on
```

For example, the following command powers up compute node 1:

```
ipmitool -I lanplus -H 10.0.40.1-U admin -P admin power on
```
4. Change the headnode hostname and IP address to the desired settings for your site, and change the headnode BMC IP address to the desired setting for your site. Refer to the *SGI Management Center Quick Start Guide*, P/N 007-5672-00x.

## Powering Up a System with the Windows HPC OS Installed

Use the following procedure to power up a system that has the Windows HPC OS installed on the headnode:

1. Verify that the power cords are connected to the rear of the system and the site power source.
2. Move the system enable switch on the rear of the enclosure to the ON position

**Note:** The system enable switch applies power to all internal trays and switches. The headnode (located in the top tray) is the only node that is configured to power-on automatically. The compute nodes do not power-on automatically; you must start them manually with the Microsoft HPC Cluster Manager or IPMI View tools. (These tools are not available until you run the `HPCGettingStarted` application.)

3. Run the `HPCGettingStarted` application from the headnode (choose **Start** -> **All Programs** -> **local Disk(C:) -> OEM Extras -> w2k8r1\_hpc\_pack\_sp1 -> HPCGettingStarted**) to install/configure the compute nodes.

**Note:** Once you have completed running the `HPCGettingStarted` application, you can use the Microsoft HPC Cluster Manager or IPMI View tools to remotely manage the compute nodes.

4. Change the headnode hostname and IP address to the desired settings for your site, and change the headnode BMC IP address to the desired setting for your site.

## Useful Information and Recommendations

- For systems that use the SGI Management Center:
  - SGI has copied the installation ISO images to the `/opt/sgi/Factory-Install` directory on the headnode. SGI recommends backing up this directory and its contents to another system or disk to enable quick recovery if the system headnode disk fails or gets corrupted. Refer to the *SGI Management Center Quick Start Guide*, P/N 007-5672-00x.
  - SGI recommends cloning the entire headnode system disk and copying it to a different system for safekeeping. Refer to the *SGI Management Center Quick Start Guide*, P/N 007-5672-00x.
  - The `/data1/sgi_host_clone.tar.gz` file includes the payloads installed on the compute nodes. SGI recommends backing this file up to another system for safekeeping. Refer to the *SGI Management Center Quick Start Guide*, P/N 007-5672-00x.
- To use `ipmitool` directly on a node (in-band), you must issue the following commands to enable use on the node:

```
# chkconfig ipmi
# /etc/init.d/ipmi start
```
- To use `ipmitool` directly on a node (in-band), enter the following command: `ipmitool <command>`

## About Non-integrated Systems

Factory-integrated systems are ready-to-use clusters. If your system was not factory-integrated, you can expect the following:

- The BMC addresses are set to the static IP addresses shown in this document.
- If the system is ordered without an operating system, you must install an operating system on the headnode before you can power on or manage the compute nodes. Once the operating system is installed, you can use the tools available with the operating system (for example, `ipmitool`) to power on and manage the compute nodes.
- If the system was ordered with the Linux OS, the operating system is installed only on the headnode. No operating system is installed on the compute nodes. Use `ipmitool` from the headnode to power on and manage the compute nodes.
- If the system was ordered with the Windows HPC OS, the operating system is installed only on the headnode. No operating system is installed on the compute nodes. Run the `HPCGettingStarted` application from the headnode (choose **Start** -> **All Programs** -> **local Disk(C:)** -> **OEM Extras** -> **w2k8r1\_hpc\_pack\_sp1** -> **HPCGettingStarted**) to install/configure the compute nodes.
- No cluster management software is installed. You must install your own cluster management software, configure the nodes, and load images on the compute nodes.

## Related Documentation

Refer to the following documents for additional information:

- *SGI Octane III Deskside Cluster Server User's Guide*, P/N 007-5633-00x
- Systems with the SGI Management Center installed:
  - *SGI Management Center Quick Start Guide*, P/N 007-5672-00x
  - *SGI Management Center Installation and Configuration*, P/N 007-5655-00x
  - *SGI Management Center System Administrator's Guide*, P/N 007-5642-00x
- Systems with Windows HPC installed:
  - *IPMI View User's Guide*
- The documents that shipped with the system (customer letter and manufacturing configuration/integration documentation)

© 2011, Silicon Graphics International Corp. All rights reserved.

Silicon Graphics, SGI, and the SGI logo are registered trademarks and ISLE and Octane are trademarks of Silicon Graphics International Corp., in the United States and/or other countries. worldwide.

Linux is a registered trademark of Linus Torvalds in several countries. Microsoft and Windows are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. All other trademarks mentioned herein are the property of their respective owners.

