

```
-A INPUT -i 10.10.10.200 -m state --state NEW -m multiport -p tcp -s  
10.10.10.0/24 -d 10.10.10.0/24 --dports 16851 -j ACCEPT
```

Example 2.5. Port 16851: modclusterd

```
-A INPUT -i 10.10.10.200 -m state --state NEW -m multiport -p tcp -s  
10.10.10.0/24 -d 10.10.10.0/24 --dports 21064 -j ACCEPT
```

Example 2.6. Port 21064: dlm

```
-A INPUT -i 10.10.10.200 -m state --state NEW -m multiport -p tcp -s  
10.10.10.0/24 -d 10.10.10.0/24 --dports 41966,41967,41968,41969 -j ACCEPT
```

Example 2.7. Ports 41966, 41967, 41968, 41969: rgmanager

```
-A INPUT -i 10.10.10.200 -m state --state NEW -m multiport -p tcp -s  
10.10.10.0/24 -d 10.10.10.0/24 --dports 50006,50008,50009 -j ACCEPT
```

Example 2.8. Ports 50006, 50008, 50009: ccsd (TCP)

```
-A INPUT -i 10.10.10.200 -m state --state NEW -m multiport -p udp -s  
10.10.10.0/24 -d 10.10.10.0/24 --dports 50007 -j ACCEPT
```

Example 2.9. Port 50007: ccsd (UDP)

2. Configuring ACPI For Use with Integrated Fence Devices

If your cluster uses integrated fence devices ([Table 2.3, “Integrated Fence Devices Supported in a Red Hat Cluster”](#)), you must configure ACPI (Advanced Configuration and Power Interface) to ensure immediate and complete fencing.

If a cluster node is configured to be fenced by an integrated fence device, disable ACPI Soft-Off for that node. Disabling ACPI Soft-Off allows an integrated fence device to turn off a node immediately and completely rather than attempting a clean shutdown (for example, `shutdown -h now`). Otherwise, if ACPI Soft-Off is enabled, an integrated fence device can take four or more seconds to turn off a node (refer to note that follows). In addition, if ACPI Soft-Off is enabled and a node panics or freezes during shutdown, an integrated fence device may not be able to turn off the node. Under those circumstances, fencing is delayed or unsuccessful. Consequently, when a node is fenced with an integrated fence device and ACPI Soft-Off is enabled, a cluster recovers slowly or requires administrative intervention to recover.



Note

The amount of time required to fence a node depends on the integrated fence device used. Some integrated fence devices perform the equivalent of pressing and holding the power button; therefore, the fence device turns off the node in four to five seconds. Other integrated fence devices perform the equivalent of pressing the power button momentarily, relying on the operating system to turn off the node; therefore, the fence device turns off the node in a time span much longer than four to five seconds.

To disable ACPI Soft-Off, use `chkconfig` management and verify that the node turns off immediately when fenced. The preferred way to disable ACPI Soft-Off is with `chkconfig` management: however, if that method is not satisfactory for your cluster, you can disable ACPI Soft-Off with one of the following alternate methods:

- Changing the BIOS setting to "instant-off" or an equivalent setting that turns off the node without delay



Note

Disabling ACPI Soft-Off with the BIOS may not be possible with some computers.

- Appending `acpi=off` to the kernel boot command line of the `/boot/grub/grub.conf` file



Important

This method completely disables ACPI; some computers do not boot correctly if ACPI is completely disabled. Use this method *only* if the other methods are not effective for your cluster.

Fence Device	Description
DRAC	Dell™ Remote Access Controller — Provides remote management, system recovery, and power control of Dell PowerEdge™ systems.
HP iLO	HP Integrated Lights Out Management Processor — Provides remote access to server console, server power control and hardware reset, and remote network booting.
IBM RSA II	IBM Remote Supervisor Adapter II — Provides remote management and power control of IBM xSeries servers.
IPMI LAN	Intelligent Platform Management Interface over LAN — Provides remote management and power control of servers via IPMI over a LAN.

Table 2.3. Integrated Fence Devices Supported in a Red Hat Cluster



Note

For the most current information about integrated fence devices supported by Red Hat Cluster Suite, contact an authorized Red Hat representative.

The following sections provide procedures for the preferred method and alternate methods of disabling ACPI Soft-Off:

- [Section 2.1, “Disabling ACPI Soft-Off with `chkconfig` Management”](#) — Preferred method
- [Section 2.2, “Disabling ACPI Soft-Off with the BIOS”](#) — First alternate method
- [Section 2.3, “Disabling ACPI Completely in the `grub.conf` File”](#) — Second alternate method

2.1. Disabling ACPI Soft-Off with `chkconfig` Management

You can use `chkconfig` management to disable ACPI Soft-Off either by removing the ACPI daemon (`acpid`) from `chkconfig` management or by turning off `acpid`.



Note

This is the preferred method of disabling ACPI Soft-Off.

Disable ACPI Soft-Off with `chkconfig` management at each cluster node as follows:

1. Run either of the following commands:
 - `chkconfig --del acpid` — This command removes `acpid` from `chkconfig` management.
 - OR —
 - `chkconfig --level 2345 acpid off` — This command turns off `acpid`.
2. Reboot the node.
3. When the cluster is configured and running, verify that the node turns off immediately when fenced.



Tip

You can fence the node with the `fence_node` command or **Conga**.

2.2. Disabling ACPI Soft-Off with the BIOS

The preferred method of disabling ACPI Soft-Off is with `chkconfig` management ([Section 2.1, “Disabling ACPI Soft-Off with `chkconfig` Management”](#)). However, if the preferred method is not effective for your cluster, follow the procedure in this section.



Note

Disabling ACPI Soft-Off with the BIOS may not be possible with some computers.

You can disable ACPI Soft-Off by configuring the BIOS of each cluster node as follows:

1. Reboot the node and start the BIOS CMOS Setup Utility program.
2. Navigate to the **Power** menu (or equivalent power management menu).
3. At the **Power** menu, set the **Soft-Off by PWR-BTTN** function (or equivalent) to **Instant-Off** (or the equivalent setting that turns off the node via the power button without delay).
*Example 2.10, “BIOS CMOS Setup Utility: **Soft-Off by PWR-BTTN** set to **Instant-Off**”* shows a **Power** menu with **ACPI Function** set to **Enabled** and **Soft-Off by PWR-BTTN** set to **Instant-Off**.

**Note**

The equivalents to **ACPI Function**, **Soft-Off by PWR-BTTN**, and **Instant-Off** may vary among computers. However, the objective of this procedure is to configure the BIOS so that the computer is turned off via the power button without delay.

4. Exit the BIOS CMOS Setup Utility program, saving the BIOS configuration.
5. When the cluster is configured and running, verify that the node turns off immediately when fenced.

**Tip**

You can fence the node with the `fence_node` command or **Conga**.

ACPI Function	[Enabled]	Item Help
ACPI Suspend Type	[S1(POS)]	
x Run VGABIOS if S3 Resume	Auto	Menu Level *
Suspend Mode	[Disabled]	
HDD Power Down	[Disabled]	
Soft-Off by PWR-BTTN	[Instant-Off]	
CPU THRM-Throttling	[50.0%]	
Wake-Up by PCI card	[Enabled]	
Power On by Ring	[Enabled]	
Wake Up On LAN	[Enabled]	
x USB KB Wake-Up From S3	Disabled	
Resume by Alarm	[Disabled]	
x Date(of Month) Alarm	0	
x Time(hh:mm:ss) Alarm	0 : 0 : 0	
POWER ON Function	[BUTTON ONLY]	
x KB Power ON Password	Enter	
x Hot Key Power ON	Ctrl-F1	

This example shows **ACPI Function** set to **Enabled**, and **Soft-Off by PWR-BTTN** set to **Instant-Off**.

Example 2.10. BIOS CMOS Setup Utility: Soft-Off by PWR-BTTN set to Instant-Off

2.3. Disabling ACPI Completely in the `grub.conf` File

The preferred method of disabling ACPI Soft-Off is with `chkconfig` management ([Section 2.1, “Disabling ACPI Soft-Off with `chkconfig` Management”](#)). If the preferred method is not effective for your cluster, you can disable ACPI Soft-Off with the BIOS power management ([Section 2.2, “Disabling ACPI Soft-Off with the BIOS”](#)). If neither of those methods is effective for your cluster, you can disable ACPI completely by appending `acpi=off` to the kernel boot command line in the `grub.conf` file.



Important

This method completely disables ACPI; some computers do not boot correctly if ACPI is completely disabled. Use this method *only* if the other methods are not effective for your cluster.

You can disable ACPI completely by editing the `grub.conf` file of each cluster node as follows:

1. Open `/boot/grub/grub.conf` with a text editor.
2. Append `acpi=off` to the kernel boot command line in `/boot/grub/grub.conf` (refer to [Example 2.11, “Kernel Boot Command Line with `acpi=off` Appended to It”](#)).
3. Reboot the node.
4. When the cluster is configured and running, verify that the node turns off immediately when fenced.



Tip

You can fence the node with the `fence_node` command or **Conga**.

```
# grub.conf generated by anaconda
#
# Note that you do not have to rerun grub after making changes to this file
# NOTICE: You have a /boot partition. This means that
#           all kernel and initrd paths are relative to /boot/, eg.
#           root (hd0,0)
#           kernel /vmlinuz-version ro root=/dev/VolGroup00/LogVol100
#           initrd /initrd-version.img
#boot=/dev/hda
default=0
timeout=5
serial --unit=0 --speed=115200
terminal --timeout=5 serial console
```

```
title Red Hat Enterprise Linux Server (2.6.18-36.el5)
    root (hd0,0)
    kernel /vmlinuz-2.6.18-36.el5 ro root=/dev/VolGroup00/LogVol100
console=ttyS0,115200n8 acpi=off
    initrd /initrd-2.6.18-36.el5.img
```

In this example, `acpi=off` has been appended to the kernel boot command line — the line starting with "kernel /vmlinuz-2.6.18-36.el5".

Example 2.11. Kernel Boot Command Line with `acpi=off` Appended to It

3. Configuring `max_luns`

If RAID storage in your cluster presents multiple LUNs (Logical Unit Numbers), each cluster node must be able to access those LUNs. To enable access to all LUNs presented, configure `max_luns` in the `/etc/modprobe.conf` file of each node as follows:

1. Open `/etc/modprobe.conf` with a text editor.
2. Append the following line to `/etc/modprobe.conf`. Set *N* to the highest numbered LUN that is presented by RAID storage.

```
options scsi_mod max_luns=N
```

For example, with the following line appended to the `/etc/modprobe.conf` file, a node can access LUNs numbered as high as 255:

```
options scsi_mod max_luns=255
```

3. Save `/etc/modprobe.conf`.
4. Run `mkinitrd` to rebuild `initrd` for the currently running kernel as follows. Set the `kernel` variable to the currently running kernel.

```
# cd /boot
# mkinitrd -f -v initrd-kernel.img kernel
```

For example, the currently running kernel in the following `mkinitrd` command is 2.6.9-34.0.2.EL: