Dell PowerEdge C System Management

A Dell Quick Start Guide

Dell PowerEdge C Servers

Table of Contents

Summary

This guide will quickly walk you through the process of bringing up a new, factory-sealed Dell™ PowerEdge™ C server, and to provide quick introductions to essential technologies along the way.

The PowerEdge C line of servers uses a baseboard management controller (BMC) to provide system monitoring and administration, including remote power-on.

The BIOS and BMC work mostly independently. The system can boot up normally even with the BMC disabled, although all advanced monitoring and out-of-band management will be unavailable.

Dell provides the following Linux command-line tools to assist managing PowerEdge C servers (all tools work with all PowerEdge C offerings). Please note that while these tools are not officially supported, they should work for all PowerEdge C servers.

*Note: **in-band** means the tool must be run on the target server. **out-of-band** means the tool can be run from anywhere on the BMC's network, and the BMC does the work—although out-of-band tools can also run on the target server if desired.*

Get these tools at: http://12.205.29.196/pec/

- **setupbios** – in-band – (no requirements, Linux)

    o Show and edit BIOS settings from command line
    o Save settings to a file; read settings from a file; scale out to many Cloud servers
    o Show only settings that were changed from defaults (to visualize)
    o Reset each setting to default
    o Set BIOS boot device order

- **bmc** – out-of-band – (Requires ipmitool, and a basic install of Perl 5.8 or higher)

    o Communicates with the always-on BMC that is on every PowerEdge-C server node
    o Agentless; driverless
    o Autodetects the remote server type
    o Manage a box remotely over the network
    o Provides convenience functions to the user
        ▪ Power a node up or down
        ▪ Probe the node number (in the chassis)
        ▪ Set the BMC ethernet to shared/dedicated

- Change GUI timeout
- Toggle blue chassis ID light
- (and many others)

- **ipmiflash** – out-of-band – (Requires a basic install of Perl 5.8 or higher)

  - Flash update the BIOS, BMC, or FCB firmware over the network via the BMC. *Note: Updating BIOS will cause a hard, immediate shutdown of the target server. Updating BMC / FCB does not require system reboot. (Requires a basic install of Perl 5.8 or higher.)*

- **ipmitool** – out-of-band

  - An open source tool, available on sourceforge. (Not provided by Dell, but heavily utilized.)

- **nagios script pack** – (upcoming, TBD): Will provide drop-in monitoring support for all PowerEdge C platforms.

- **PXE image (provisioning)**

  - Simple, small CentOS image that is preconfigured for PE-C
  - Driven by simple bash scripts, dynamically pulled down from FTP server at boot
  - Flexible; multiple uses triggered by simple flags
    - Inventorying (Firmware levels; server identity; almost anything)
    - Pull down scripts and tools from FTP server at boot-time
    - Pushing collected information back up to the FTP server
    - Provisioning and deployment
    - BIOS, BMC, FCB, LSI firmware updates (if required)
    - Decommissioning and repurposing

In addition to these tools, to manage storage the following is recommended (you must get these tools directly from LSI's website; search support for the LSI 9260-8i):

- **MSM (MegaRAID Storage Management)** – in-band

  - GUI-based. Monitor and manage storage; includes the most common settings and options.

- **MegaCLI** – in-band

  - All-in-one command line tool to manage everything regarding storage: Create and manage logical volumes; probe for status and errors; drop, replace, and reconstruct drives and arrays. Very flexible, comprehensive, and useful in scripting mass deployments.

Bare Metal Setup (Prepare the BMC for Operation)

1. Put the BMC on the network.

    a. If you are using "shared" Ethernet cabling mode, the cable goes into LOM1 (LAN-on-Motherboard 1). Shared mode is the default; it reduces Ethernet cabling and will appear on the network as two separate NICs (presenting two distinct MAC addresses): LOM1 and the BMC.

    b. If you are using "dedicated" Ethernet cabling mode, the cable goes into the wrench-logoed RJ-45. This allows for physical separation of the management network.

2. Plug in AC power. It should take about 1 minute for the BMC to boot. It will come up automatically once AC power is applied.

3. Get the BMC MAC address from BIOS settings:

    ```
    PowerEdge C6100/C6145:  Server | IPMI configuration
    All other models:       Server | Set BMC LAN configuration
    ```

4. Verify the BMC has acquired an IP address (if using DHCP). Sometimes there can be a delay in this process. In the worst case, there should be at most a 10-15 minute delay. Switching from "shared" to "dedicated" and back again (no need to exit BIOS) can force this.

Prepare the Management Server to Interact with the BMC

1. Verify that the management (client) computer can ping your server's BMC. (It will be easiest to use a Linux instance for this purpose.)

2. To most easily interact with the PE-C's BMC: Get and build IPMItool locally, if you don't already have it. An older existing version will likely be fine. The Dell-forked versions are not necessary (you may find reference to them on the Internet). Get it from:

http://ipmitool.sourceforge.net/

Build IPMItool with the following commands:

```
./configure
make
make install
```

3. Get the Dell PowerEdge C System Management Pack—it is small, and contains tools to manage the servers.

4. The BMC should now be offering a Web GUI you can access with your browser.

http://my.bmc.ip.addr

Managing the Server Via a Remote Console

PowerEdge C servers offer, as standard, many options of remotely managing the server. All of the following methods will effectively allow configuration of BIOS setup options and full OS installation.

- **iKVM**: Virtual KVM over the network, provided by the BMC. Access via "remote control" inside the BMC Web GUI. Requires Oracle Java on the client machine.
- **SOL**: Serial Over LAN. Fully virtualized serial interface—handled by the BMC. Requires IPMItool.
- **Serial port**: As above, but over a real physical serial port, so it can be redirected via a concentrator.

Notes Regarding Serial Console

Using the serial port can be an efficient means of managing your server, but it does require some work to set up.

Parameters must be configured into the Linux bootloader and provided as a kernel command line parameter for a smooth handoff. (Otherwise, output will just cease). You must ensure that the serial port setup (data rate, stop bits, etc.) matches between:

- The console itself: a terminal emulator such as PuTTY, or an instance of "ipmitool sol activate"
- The server's BIOS setup
- The Linux bootloader (likely "grub")
- The Linux kernel itself (via command line parameter "console=")
- The "getty" instances referenced in /etc/inittab

Expect to have something like this appended to the Linux kernel command line (via the bootloader):

```
console=tty0 console=ttyS1,19200n1
```

Note that multiple "console=" entries are allowed—the last one listed sees the startup messages (stdin/stdout/stderr of /sbin/init).

Using iKVM

iKVM offers a virtual console for your PowerEdge C server over the network. Here are the steps to try it out:

- Install Java on your client machine.
- Point a Web browser at the BMC's IP address (https://my.bmc.ip.addr).
- Log in (defaults are root / root).
- Click "console" tab.
- Click "launch kvm" or "launch virtual console."

You can manage the BIOS options, install an OS interactively, and so on. Virtual media can also be mounted—media local to your client machine will appear as media local to the server, across the network.

It is easiest to enable "full keyboard," so that all keys are captured and passed to the server:

```
PowerEdge C1100/C2100:  Tools | Session Options | Pass all
keystrokes to target

All other models:  Keyboard | Full Keyboard
```

Pasting into the iKVM window is available for C5125 / C5220, but **copying** from the window is not possible. (This is a fundamental limitation: Remotely, keystrokes and mouse movements are being emulated by the BMC at the hardware level. This tradeoff allows the use of any, or no, OS).

Using IPMI

You can do most system management and monitoring via IPMI, from anywhere on the network. If you invoke using IPMI v2.0 (by specifying "`-Ilanplus`"), all IPMI communications are encrypted and secure.

The general structure of commands is:

```
ipmitool [CONNECTION INFO] [COMMANDS]
```

Where the connection info will be something like (default user is root / root):

```
-Ilanplus -H192.168.0.5 -Uroot -Proot
```

And the command will be something like (to print the System Event Log):

```
sel list
```

So to put it all together, here's an example:

```
ipmitool -Ilanplus -H192.168.0.5 -Uroot -Proot sel list
```

Commands are hierarchical. Sending "sel help" will show all available SEL commands. Here are some commands that are particularly useful:

- Show all top-level commands

  ```
  help
  ```

- Remote power on/off commands

  ```
  chassis help
  chassis power on
  chassis status
  ```

- Show BMC LAN connection info

  ```
  lan print 1
  ```

- System Event Log (SEL)

```
sel
sel list
```

- Sensor Data Records (SDR)

```
sdr list
sdr elist
sdr list -v
```

- Management Console—info about the BMC itself.

```
mc info
mc reset cold          # reboot the BMC
```

- FRU info (what's in the board)

```
fru print
```

- Serial over LAN (includes complete preboot management; setting up BIOS settings)

```
sol activate
```

Notes on SOL: Once started, press **Return** *followed by the following sequences:*

```
~?          Help
~.          Exit SOL mode
```

If SOL mode doesn't work by default, go into BIOS setup:

```
Server | Remote Access
```

Change the port to COM2, and set the speed appropriately.

Using the "setupbios" Tool

The setupbios tool must be run on the target machine from a Linux command prompt.

Configuring BIOS Settings

The tool can read and write BIOS settings in several ways.

To display settings:

```
# ./setupbios settings show
```

To display in an alternate format, or save the settings (this is a good way to capture settings on one system, to be applied to another):

```
# ./setupbios settings save
# ./setupbios settings save > mysettings.txt
```

To display only settings that differ from defaults:

```
# ./setupbios settings save_diffs
```

To read and apply settings from a saved file:

```
# ./setupbios settings readfile mysettings.txt
```

Configuring Boot Order

Currently, the setupbios tool allows the user to change the first eight BIOS boot order slots. This tool does not attempt to protect the user from invalid input. To read or set BIOS device boot order:

```
# ./setupbios bootorder [platform] list|get|set [index] [device_num]
```

An example of this in action, setting the first five slots on a PowerEdge C6100:

```
# ./setupbios bootorder c6100 set 1 4
# ./setupbios bootorder c6100 set 2 1
# ./setupbios bootorder c6100 set 3 2
# ./setupbios bootorder c6100 set 4 0
```

```
# ./setupbios bootorder c6100 set 5 3
# ./setupbios bootorder c6100 list
Boot #1:  device id: 4
Boot #2:  device id: 1
Boot #3:  device id: 2
Boot #4:  device id: 0
Boot #5:  device id: 3
```

Using the "bmc" Tool

The "bmc" tool abstracts some of the more complicated IPMItool commands away and provides a simple, Linux-like interface. It works with all PowerEdge C platforms and attempts to provide the most commonly used commands in an easy format, while providing access to all architected extended configuration attributes.

Commonly used commands:

- `version` – Probe a BMC, return its firmware version.

- `nic_mode` – Probe or set shared/dedicated Ethernet mode.

- `platform` – Probe the target BMC and report back the platform type.

- `reset` – Cause a BMC reboot.

- `node_number` – In case of shared infrastructure chassis, report which node we are.

- `bmc_hostname` – Get or set the BMC's hostname.

- `gui_timeout` – Get or set the BMC's Web GUI timeout. 0 sometimes means unlimited; it may also mean "instant timeout."

- `id_led` – Set the blue ID light on/off.

- Extended IPMI attributes are available via the "bmc attr" command. It is used as follows:

```
attr  list|get|set  [attr_name]  [attr_value]
```

- "`attr list`" – will show all defined attributes, but not all defined extended attributes are available on all platforms.

Using the "ipmiflash" Tool

Find the firmware to be flashed from www.support.dell.com. Place the firmware image in a location that is network-accessible by TFTP, FTP (user/password is OK), or HTTP. Then issue the command (can be from a remote machine):

```
# ./ipmiflash [connection Info] bios|bmc|fcb ftp://user:pass@host:port/path/to/ROM
```

Some validation checks will be automatically performed prior to updating.

SNMP, Dell OpenManage ITA, and Other Consoles

PowerEdge and PowerEdge C boxes do not share a common lineage. *Dell PowerEdge MIBs will not work.*

Dell OMSA can't be run on the PowerEdge C series because OM (ITA, Bandera) has no explicit support for PowerEdge C. You can use ITA's IPMI to probe the PowerEdge C server from ITA or Bandera, but you only get very minimal information because the detection methods used are custom to PowerEdge. PowerEdge C boxes are more focused on using generic, industry-standard methods to do health monitoring. As such, third-party tools have been shown to successfully probe and manage PowerEdge C servers.

PowerEdge C platforms currently do not offer SNMP with a MIB. PowerEdge C does offer limited SNMP, throwing generic traps that contain system event log (SEL) messages, but no included MIB. System health monitoring should instead be done by actively probing the sensors via IPMI.

The BMC aggregates sensor values that can be probed via IPMI ("`ipmitool sdr list`" or "`ipmitool sensor`"). RAID adapters can be probed directly via third party tools (LSI's MegaCli) to look for faults.

The BMC has a System Event Log (SEL). Most significant health problems related to the system will show up in this log. Storage must be separately monitored.

Health Monitoring Strategy

Dell recommends using Nagios to monitor PowerEdge C boxes. Configure Nagios to:

- Using ipmitool, poll the BMC's error log (SEL) for errors; record these; clear the SEL if full.
- Using ipmitool, poll and parse the BMC's SDR list to monitor sensors, and look for ones that are out of spec:

  - Temperature
  - Voltage
  - Fan RPM
  - NMI flags
- Poll and parse MegaCLI inside the OS, looking for any degraded arrays or other errors.
- Poll relevant services/servers running on the OS, verifying that they are alive and responsive.

Provisioning the Server

Dell can provide, at no charge, a bootable PXE image available that can facilitate provisioning, inventorying, configuring, and firmware updates for your PowerEdge C server. This bootable image can also be easily modified to your needs, or it can serve as a guide for creating your own.

Configuring Storage

In most cases (LSI adapters, Dell PERC adapters), storage is provisioned with a third-party tool: "megacli." One exception is the LSI2008 card; this uses "lsiutil." Get and install these tools directly from LSI.

Included here is an example of drive provisioning, taken from the Dell PowerEdge C provisioning PXE image. These storage adapters organize drives by:

- Adapter (numbered 0, 1, 2, … sequentially)
- Enclosure
- Slot [in an enclosure]

Example follows:

```
#
# Clear all current configuration
#
$MEGACLI –CfgClr –aALL
$MEGACLI –CfgForeign –Clear –aALL


#
# Configure Raid 0
# $MEGACLI –CfgLdAdd –r0 [${LSI_ENC_ID}:${LSI_SLOTS[0]}] –a0
#


#
# Configure Raid 5
#
#$MEGACLI –CfgLdAdd –r5     [${LSI_ENC_ID}:${LSI_SLOTS[0]}, \
#                            ${LSI_ENC_ID}:${LSI_SLOTS[1]}, \
#                            ${LSI_ENC_ID}:${LSI_SLOTS[2]}, \
#                            ${LSI_ENC_ID}:${LSI_SLOTS[3]}] \
#                       –a${LSI_ADPNUM}


#
# Configure Raid 10
#
$MEGACLI –CfgSpanAdd –r10 \
     –Array0[${LSI_ENC_ID}:${LSI_SLOTS[0]},${LSI_ENC_ID}:${LSI_SLOTS[1]}] \
     –Array1[${LSI_ENC_ID}:${LSI_SLOTS[2]},${LSI_ENC_ID}:${LSI_SLOTS[3]}] \
          –a${LSI_ADPNUM}
```

The adapter information, array configuration, and logical drive / physical drive information can later be pulled:

```
# Show adapter configuration
$MEGACLI –AdpAllInfo –aAll

# Show enclosure contents summaries
$MEGACLI –encinfo –aAll

# Show all LD configuration, and the drives that comprise an LD
$MEGACLI –CfgDsply –aAll

# Display Logical and Physical drives on a specified adapter
$MEGACLI –LdPdInfo –aAll
```

There are a number of great quick reference guides with recipes on the Web to help you get started using this tool.

Booting from 3TB Drives

Refer to the following URL for an extended description:

http://www.delltechcenter.com/page/HDD+Support+for+2.5TB,+3TB+Drives+and+Beyond

Large drives present a new issue: Legacy OS and BIOS only support "MBR" partitions, but those are limited to 2.2 TB at maximum. UEFI can boot from "GPT" partitions, which support extremely large amounts of storage.

The workaround with legacy BIOS is to use the adapter to carve off an MBR partition that is 2.2 TB (or less) on the boot drive. This can be configured in the "WebBIOS" (during system boot-up) or via megacli (check help for options on "-CfgLdAdd").

Additionally, Linux has some support for multiple MBR tables, which would be another way to carve up a large drive. This is not present in Windows.

To Learn More

For more information on the Dell PowerEdge C servers, visit:
www.dell.com/poweredgec