



Installing Red Hat Enterprise Linux Advanced Server 3 on IBM @server BladeCenter JS20

October 19, 2004

Written By:
Erik Salander
IBM Corporation
salander@us.ibm.com

Contributor:
Kevin McCombs
IBM Corporation
mccombsk@us.ibm.com



I. Introduction.....	3
Prerequisites.....	3
IP address scheme	4
Using the Management Module's Web interface.....	4
Remote access to the JS20.....	4
Using a Serial Over LAN (SOL) session.....	5
II. Manual installation from CD.....	5
Setup to install from CD	5
Start the installation.....	6
Using the Installation Program.....	6
III. Manual installation over network	8
Setup of the install server.....	9
DHCP configuration.....	9
TFTP configuration	10
NFS configuration.....	11
Building the installation tree	11
Setup to install over network	12
Start the installation and define installation settings	13
Using the Installation Program.....	13
IV. Auto-installation over network via OpenFirmware	15
Setup of the install server.....	16
Create the kickstart configuration file	16
Install from OpenFirmware	16
V. Auto-installation over network using CSM.....	18
CSM management server (MS) setup.....	18
CSM software package from Website	19
CSM software package from CDs	20
Accept the product license	20
Define the Nodes to the MS	22
CSM Red Hat (kickstart) installation setup.....	26
Install the OS to the nodes	27
Monitoring with CSM - optional.....	29
VI. Red Hat Network.....	30
RHN components.....	31
RHN service levels.....	31
RHN Update service overview.....	31
VII. Troubleshooting	32
Serial Over LAN problems.....	32
DHCP problems.....	32
TFTP problems	33
CSM problems.....	33
VIII. References.....	34



I. Introduction

This white paper explains how to install Red Hat Enterprise Linux® Advanced Server 3 for POWER™ Update 3 (RHEL AS 3 Update 3) onto an IBM eServer® BladeCenter™ JS20 using a variety of install procedures. The install procedures, their incremental differences and their typical use case are as follows:

- Manual installation from CD
 - no install server, step through the Installation Program
 - use when there are few systems, infrequent installs
- Manual installation over network
 - create install server, step through the Installation Program
 - use when installing more systems with varying configurations, or short-lived configurations
- Auto-installation over network
 - create install server and kickstart file
 - use when frequently installing many systems with a similar configuration or categories of configurations
- Auto-installation over network using Cluster Systems Management (CSM)
 - CSM sets up the install server, kickstart files and coordinates the installation
 - use when a full function, yet easy-to-use distributed system management solution is necessary for a cluster of nodes

It's also possible to use a combination of the install procedures listed above. For instance, even if there's an install server, containing an install tree and kickstart files. A manual installation over the network could be carried out, only taking advantage of the install tree.

The remainder of this section describes prerequisites, network information and procedures that will be reused in multiple sections of the document.

Prerequisites

This document assumes the Serial Over LAN (SOL) feature is already setup on the BladeCenter JS20. The JS20 has no direct serial connection for a monitor, mouse or keyboard. Therefore, to enable communication between a blade server and these devices, you must configure the Serial Over LAN (SOL) function of the blade server. To find out more about SOL setup, refer to the *Serial over LAN Setup Guide – IBM eServer™ BladeCenter and BladeCenter T* document.

Before you get started

When installing Red Hat Enterprise Linux AS 3 Update 3 on the BladeCenter JS20 2.2 GHz, ensure the following settings are made.

eth1 is static with ONBOOT=yes
eth0 is dhcp with ONBOOT=no

SMS is not enabled on the JS20 2.2 GHz. You should use Web ASM to change the boot order.



Clients can get the latest updates and service packs from Red Hat Linux by registering on their maintenance Web sites and downloading the latest versions, updates or erratas.

IP address scheme

This is the IP address scheme used throughout this document. The systems involved consist of a couple JS20 blades and an install server. The two blades in the BladeCenter are connected to an install server on a private network. The IP addresses are as follows:

Install Server:

- IP address is 192.168.70.50/255.255.255.0

BladeCenter components:

- management module external is static IP at 192.168.70.125/255.255.255.0
- management module to blades is static IP at 192.168.70.126/255.255.255.0
- management module to IO modules is static IP at 192.168.70.127/255.255.255.0
- Ethernet switch module 1 is static IP at 192.168.70.130/255.255.255.0
- Ethernet switch module 2 is static IP at 192.168.70.131/255.255.255.0

Blade1:

- Eth0 is disabled, to accommodate only SOL usage
- Eth1 is dynamic IP at 192.168.70.100/255.255.255.0

Blade2:

- Eth0 is disabled, to accommodate only SOL usage
- Eth1 is dynamic IP at 192.168.70.101/255.255.255.0

Gateway:

- IP address is 192.168.70.1

Using the Management Module's Web interface

You'll often need to use the BladeCenter's Management Module Web interface, since the MM provides the central point of control for the BladeCenter components. To log into the MM Web interface, start a browser and point it to the IP address or name of the MM (e.g. <http://192.168.70.125>). Enter the userid/password, the default is USERID/PASSWORD (note: all upper case, with a zero). This is usually done from a ThinkPad®, or whatever system happens to be available. Select a reasonable timeout value then click on Continue.

Remote access to the JS20

For some Linux distributions we use the Serial Over LAN (SOL) session to perform text-based installs and Virtual Network Computing (VNC) connections to perform GUI-based installs. At the moment, there exist some problems using VNC on JS20's with RHEL. This whitepaper will only utilize SOL sessions to drive/view the install process.



Using a Serial Over LAN (SOL) session

To view the serial console output, use the Serial Over LAN feature of the BladeCenter. The *Serial over LAN Setup Guide – IBM eServer BladeCenter* describes how to setup the components needed for SOL functionality.

An SOL session accesses the JS20 via the MM. To start a SOL session from an attached system (eg. a ThinkPad, etc...), telnet to the MM, enter "telnet 192.168.70.125". When prompted, provide the userid/password, the default is USERID/PASSW0RD (note: all upper case, with a zero). To bring up a console for a blade, enter this command at the telnet prompt: "console -T system:blade[x]", where x is a number 1-14, indicating which blade in the chassis to select. The SOL session buffers up to 8KB of data, so when the session is started, any buffered data will scroll past.

To set the "environment", so all commands are directed to specific blade, enter: "env -T system:blade[x]", where x is a number 1-14. Then one just needs to enter: "console", to start an SOL session on blade x.

To terminate an SOL session, press ESC, then shift-9 (a left parenthesis).

The Windows® telnet client is not good at displaying the installation program screens. The free telnet client PUTTY seems to work much better. PUTTY can be downloaded from: <http://www.chiark.greenend.org.uk/~sgtatham/putty/download.html>

There are a few recommended changes to the standard PUTTY configuration that will make the installation program screens more useable. Start PUTTY and make these changes:

- click on Window, change the rows/columns from 24/80 (default) to 40/100
- click on Appearance, change the Font to Courier New, with a Font size of 10

Create all your PUTTY telnet configurations/sessions with those changes.

II. Manual installation from CD

This section describes the process of installation from the RHEL CDs. This install option is most useful when a small number of systems need to be infrequently installed. These are the steps necessary to perform this type of installation:

Setup to install from CD-ROM

Start the MM Web interface, as described in *Section I, Using the Management Module's Web interface*. Under "Blade Tasks", select "Configuration", then "Boot Sequence". Select the server to be installed. Insert "Hard drive 0" first in the startup sequence and CD-ROM second into the boot sequence for the server and click on "Save". This assumes that Hard drive 0 does not have a bootable system installed. If it does, insert CD-ROM first in the boot sequence. In this case, after installing to Hard drive 0, you'll need to come back and put Hard drive 0 first and CD-ROM second, in the boot sequence. Note the checkbox for changing all the boot sequence on all the blades at once.



Give the target install blade sole access to the CD-ROM drive by pressing the CD button on the top/front of the blade. This can also be done remotely in the MM Web interface, under “Blade Tasks”, select “Remote Control”, then “Start Remote Control” and select the blade from the “Change media tray owner” pulldown. Insert the RHEL AS 3 Update 3 CD #1 into the media tray’s CD-ROM drive.

Start the installation

To reboot the server, and thereby kick off the installation process, under “Blade Tasks” on the left hand side of the MM Web interface, select “Power/Restart”. Check the box for the blade you wish to start installing and power it on by selecting “Power On Blade” or “Restart Blade”.

Start an SOL session, as described in *Section 1, Using a Serial Over LAN session*, in order to control the installation. On the SOL console, you should see the progress codes scrolling past while the JS20 initializes.

When yaboot presents the “boot:” prompt, either press Enter or let it timeout. In either case, it will boot the default boot entry (which was just assigned to the CD-ROM).

Using the Installation Program

Next, you should see the RHEL AS 3 Update 3 Welcome screen. Read over the Welcome screen and press “OK” to continue. Normally, the next two screens would be about keyboard and mouse configuration. For a JS20, those screens are not displayed.

On the next screen you will be presented the Disk Partitioning Setup screen. You can choose between automatic partitioning and manual partitioning. If you’re not comfortable manually partitioning your system, it’s recommended you select automatic partitioning. Automatic partitioning allows Linux and/or non-Linux partitions to be selectively retained or removed. For more information on disk partitioning, read *Appendix C, An Introduction to Disk Partitions in the Installation Guide for the IBM pSeries® Architecture*. Some standard partitioning guidelines to follow when manually partitioning:

- A swap partition (at least 256 MB). Swap partitions are used to support virtual memory. In other words, data is written to a swap partition when there is not enough RAM to store the data your system is processing. If you are unsure about what size swap partition to create, make it twice the amount of RAM on your machine, but no more than 2048MB (or 2GB). It must be of type swap.

For example, if you have 1GB of RAM or less, your swap partition should be at least equal to the amount of RAM on your system, up to two times the RAM. For more than 1GB of RAM, 2GB of swap is recommended. Creating a large swap space partition can be especially helpful if you plan to upgrade your RAM at a later time.

- A PReP (PowerPC Reference Platform) boot partition on the first partition of the hard drive - the PReP boot partition contains the operating system kernel or the yaboot boot loader (which allows pSeries systems to boot Red Hat Enterprise Linux). On pSeries, unless you plan to boot from a floppy or network source, you must have a PReP boot partition to boot



Red Hat Enterprise Linux. The PReP boot partition should be between 4-8 MB for pSeries systems, not to exceed 10MB.

- A /boot/ partition (100MB) - the partition mounted on /boot/ contains the operating system kernel (which allows your system to boot Red Hat Enterprise Linux), along with files used during the bootstrap process. Due to the limitations of most PC firmware, creating a small partition to hold these is a good idea. For most users, a 100MB boot partition is sufficient.

Caution

If you have a RAID card, be aware that some firmware do not support booting from the RAID card. In cases such as these, the /boot/ partition must be created on a partition outside of the RAID array, such as on a separate hard drive.

- A root partition (350MB - 5.0GB) - this is where "/" (the root directory) is located. In this setup, all files (except those stored in /boot) are on the root partition. A 350MB partition allows you to install a minimal installation, while a 5.0GB root partition lets you perform a full installation, choosing all package groups.
- A /var/ partition (3.0GB or larger) - the /var/ partition is where variable data files are written. This includes spool directories and files, administrative and logging data, and transient and temporary files. Updates that are applied to Red Hat Enterprise Linux are written to /var/spool/update/ and require approximately 1.0GB of disk space available when updating your system.

After completing the disk partitioning step, the next installation program screen will allow you to set your Network Configuration for eth0. Disable the eth0 network device by deselecting "Activate on boot". Click on OK and then configure eth1. For example, blade1's configuration information is: IP Address of 192.168.70.100, and a Netmask of 255.255.255.0. Click on OK. Next, some Miscellaneous Network Settings are shown. For example, blade1's miscellaneous settings would be: Gateway of 192.168.70.1 and any Name Servers the organization has. Click OK to continue. The network configuration can be changed after the installation by running the "redhat-config-network" command.

The last piece of networking information is shown on the Hostname Configuration screen. Select to manually provide a hostname, for blade1 this would be blade1.bclab.com. Click OK to continue.

Next is the Firewall Configuration screen. We'll select "No firewall" since the servers are usually connected to a router providing firewall services. Click on OK. The firewall configuration can be changed after the installation by running the "redhat-config-securitylevel" command.

Then the Language Support Selection screen is shown. Choose any additional languages to install and support on your system. Click OK when finished.

The next screen will be the Timezone Configuration screen. Select the appropriate timezone for the server. The timezone configuration can be changed after the installation by running the "redhat-config-date" command. Click OK to continue.

Next, you'll be given the option to set the root password. The root password can be changed after the installation by running the "redhat-config-rootpassword" command. Click OK when done.



The last installation program screen is the Package Defaults screen, which allows you to select the software packages to be installed. You can accept the recommended package selection or customize the software selection. To accept the recommended selection, simply click on OK. To customize, select Customized software selection, and click on OK. When finished customizing, click on OK. Now, the installation of RHEL will begin. The installation takes a variable amount of time depending on the quantity and size of the packages selected.

While the installation is proceeding, the Package Installation screen will show the current package being installed and a running total of the installation packages and time estimate.

When the installation completes, you'll be shown the Complete screen. Prior to rebooting you'll need to set the boot sequence now that installation is complete. Start the MM Web interface, as described in *Section I, Using the Management Module's Web interface*. Under "Blade Tasks", select "Configuration", then "Boot Sequence". Select the server to modify. Insert "Hard drive 0" first in the startup sequence for the server(s) just installed and click on "Save".

Now, click on OK on the Complete screen. The system will reboot using the new installation.

When yaboot presents the "boot:" prompt, either press Enter or let it timeout. In either case it will boot the default boot entry. After one or more screens of messages scroll by, a "login:" prompt will be shown. Login as the root user and supply the root password provided during the configuration. The installation is now complete.

III. Manual installation over network

This install procedure is best used when installing many systems with varying configuration or short-lived configurations. That is, auto-installations aren't of value so we'll step through the Installation Program screens for each install. This install procedure requires:

- Setup of the install server
 1. DHCP configuration
 2. TFTP configuration
 3. NFS configuration
 4. Building the installation tree
- Setup to install over network
- Start the installation and define installation settings
- Using the Installation Program



Setup of the install server

In general, any system or combination of systems that provide DHCP, TFTP, NFS and the install tree satisfy the requirements for a network-based install server. In verifying these RHEL install procedures, we used an install server running RHEL AS 3 Update 3.

DHCP configuration

First, the install server needs to be setup as a BOOTP/DHCP server. The DHCP server package that comes with RHEL is capable of handling BOOTP and DHCP requests. Create an `/etc/dhcpd.conf` file, here's an example.

```
# dhcpd.conf

# option definitions common to all supported networks...
allow bootp;
allow booting;
always-reply-rfc1048 true;

default-lease-time 600;
max-lease-time 7200;
ddns-update-style none;
authoritative;
log-facility local7;

subnet 192.168.70.0 netmask 255.255.255.0 {
  group {
    next-server 192.168.70.50;
# filename "/tftpboot/install";
    filename "install";
    host blade1 {
      hardware ethernet 00:0d:60:1e:0e:8d;
      fixed-address blade1;
    }
    host blade2 {
      hardware ethernet 00:0d:60:1e:0e:75;
      fixed-address blade2;
    }
  }
}

# end of dhcpd.conf
```

To obtain the MAC addresses (ie. hardware ethernet entries) to be specified in the `dhcpd.conf` file, start the MM Web interface as described in *Section 1, Using the Management Module's Web interface*. Under "Monitors", select "Hardware VPD", then scroll down to "BladeCenter Server MAC Addresses". Use the "MAC Address 2" column. This corresponds to eth1, which will allow SOL to function on eth0.

Some issues to note about `dhcpd.conf`:



- Most tftp daemons change the root directory to something different than / (using the chroot command), usually /tftpboot. This is accomplished by the usage of the `-s` startup option to the tftp daemon. You can verify this by checking the "server_args" line in `/etc/xinetd.d/tftp`.
 - a. If the `-s` option is not specified as an argument, the full path needs to be specified for the "filename" in `dhcpd.conf`.
 - b. If the `-s` option is specified as an argument, only the relative path needs to be specified for the "filename" in `dhcpd.conf`.
- If you want your DHCP server to respond to DHCP requests from a certain interface, this can be handled in `/etc/sysconfig/dhcpd`. For example, coding:

```
DHCPD_INTERFACE="eth1"
```

```
..in /etc/sysconfig/dhcpd will force the DHCP server to only listen on eth1.
```

- The "next-server" and "filename" parameters inform the servers how to continue with the next step in the process, TFTP'ing the install kernel from the install server to the server(s).

To restart the DHCP daemon (`dhcpd`), run:

```
# service dhcpd restart
```

For hostname resolution, we'll just setup `/etc/hosts`, here's an example:

```
# hosts
```

```
192.168.70.50      ms.bclab.com      ms
192.168.70.125   mm.bclab.com      mm
192.168.70.126   mm-int.bclab.com  mm-int
192.168.70.130   esm1.bclab.com    esm1
192.168.70.131   esm2.bclab.com    esm2
```

```
192.168.70.100   blade1.bclab.com  blade1
192.168.70.101   blade2.bclab.com  blade2
```

```
# end of hosts
```

TFTP configuration

Next we need to setup the install server as a TFTP server. To do this, create a configuration file `/etc/xinetd.d/tftp`:

```
# description: The tftp server serves files using the trivial file transfer
# protocol. The tftp protocol is often used to boot diskless
# workstations, download configuration files to network-aware printers,
# and to start the installation process for some operating systems.
service tftp
{
    socket_type      = dgram
    protocol         = udp
    wait             = yes
    user             = root
    server           = /usr/sbin/in.tftpd
```



```
server_args = -vvv -r blksize -p -u tftpd -s /tftpboot
}
```

Note, the absence of the “disable = yes” line indicates to activate TFTP. The TFTP daemon process runs under xinetd, to activate TFTP capability, run:

```
# service xinetd restart
```

Then we need to copy the /images/netboot.img file from the RHEL AS 3 Update CD #1 to the /tftpboot directory.

Note, make sure the files in /tftpboot have the correct owner and group IDs assigned. Since the tftpd process usually does not run as the root user. Check the “server_args” line in /etc/xinetd.d/tftp. In our example, “-u tftpd” is specified. The user and group IDs of the files in /tftpboot must match the information in the system permission database for the user “tftpd”. To match the example, run this:

```
# chown tftpd:tftpd /tftpboot/netboot.img
```

NFS configuration

Next, we need to setup the install server as an NFS server. Configure NFS to export the /install-tree directory by adding it to /etc/exports. For example:

```
/install-tree * (ro,root_squash, sync)
```

To check if the NFS service is already started, run:

```
# service nfs status
```

If the service is started, you can reload the configuration file with this command:

```
# service nfs reload
```

To start the NFS service, run:

```
# service nfs start
```

Building the installation tree

Next, we need to setup the installation tree on the install server. When building the installation tree, the variables to consider are:

- Will it be exported using NFS or HTTP or FTP?
- Will it be a flat installation tree or a multiple source installation tree?

In this example, we'll export using NFS and create a flat installation tree. The installation tree will be built in the /install-tree/rhelas3u3 directory. Note, the /install-tree/rhelas3/QU3 directory is used to accommodate other distributions (ie. SUSE LINUX Enterprise Server 9 in the /install-tree/sles9). If you're only installing RHEL, you can shorten the directory path in the examples that follow. These are the steps to build this installation tree.



Copy the content of the RHEL AS 3 Update 3 CDs to disk. Be sure to specify the “periods” in the commands that follow.

```
# mkdir -p /install-tree/rhelas3/QU3
# cd /install-tree/rhelas3/QU3
# umount /mnt/cdrom
```

insert RHEL AS 3 Update 3 CD #1

```
# mount /dev/cdrom /mnt/cdrom
# cp -va /mnt/cdrom/RedHat .
# umount /mnt/cdrom
```

insert RHEL AS 3 Update 3 CD #2

```
# mount /dev/cdrom /mnt/cdrom
# cp -va /mnt/cdrom/RedHat .
# umount /mnt/cdrom
```

insert RHEL AS 3 Update 3 CD #3

```
# mount /dev/cdrom /mnt/cdrom
# cp -a /mnt/cdrom/RedHat .
# umount /mnt/cdrom
```

insert RHEL AS 3 Update 3 CD #4

```
# mount /dev/cdrom /mnt/cdrom
# cp -a /mnt/cdrom/RedHat .
# umount /mnt/cdrom
```

That completes the creation of the installation tree.

Setup to install over network

Start the MM Web interface as described in *Section I, Using the Management Module’s Web interface*. Under “Blade Tasks”, select “Configuration”, then “Boot Sequence”. Select “Network – BOOTP” from the pulldown for the first device in the boot sequence, and click on Save. Note the checkbox for changing all the blades startup sequence at once. Since this sends BOOTP requests out both Ethernet ports of the JS20, any open SOL sessions (which are on eth0) will be interrupted. However, the SOL sessions can be resumed by restarting them.



Start the installation and define installation settings

To reboot the server, and thereby kick off the installation process, under “Blade Tasks” on the left hand side of the MM Web interface, select “Power/Restart”. Check the box for the blade you wish to start installing and power it on by selecting “Power On Blade” or “Restart Blade”.

Start an SOL session as described in *Section I, Using a Serial Over LAN session*, in order to control the installation. On the SOL console, you should see the progress codes scrolling past while the JS20 initializes. If the SOL session is started soon enough, BOOTP requests will interrupt the SOL session and it will be terminated. If your SOL session returns to the MM telnet prompt, just restart the SOL session with the blade. If you get the message “SOL is not ready”, keep trying to restart the SOL session. It takes about 30 seconds until the SOL session can be restarted.

A series of screens will be shown that are used to gather installation settings.

- Choose a Language – select your language
- Installation Method – select “NFS image”
- Networking Device – select “eth1”
- Configure TCP/IP – for example, if installing blade1, specify static IP values of:
 - IP Address = 192.168.70.100
 - Netmask = 255.255.255.0
 - Default Gateway (IP) = 192.168.70.1
 - Primary Name server
- NFS Setup – specify the info for our install server:
 - NFS server name = 192.168.70.50
 - Red Hat directory = /install-tree/rhelas3/QU3

Using the Installation Program

Next, you should see the RHEL AS 3 Update 3 Welcome screen. Read over the Welcome screen and press “OK” to continue. Normally, the next two screens would be about keyboard and mouse configuration. For a JS20, those screens are not displayed.

On the next screen you will be presented the Disk Partitioning Setup screen. You can choose between automatic partitioning and manual partitioning. If you’re not comfortable manually partitioning your system, it’s recommended you select automatic partitioning. Automatic partitioning allows Linux and/or non-Linux partitions to be selectively retained or removed. For more information on disk partitioning, read *Appendix C, An Introduction to Disk Partitions in the Installation Guide for the IBM pSeries Architecture*. Some standard partitioning guidelines to follow when manually partitioning:

- A swap partition (at least 256MB). Swap partitions are used to support virtual memory. In other words, data is written to a swap partition when there is not enough RAM to store the data your system is processing. If you are unsure about what size swap partition to create, make it twice the amount of RAM on your machine, but no more than 2048MB (or 2GB). It must be of type swap.

For example, if you have 1GB of RAM or less, your swap partition should be at least equal to



the amount of RAM on your system, up to two times the RAM. For more than 1GB of RAM, 2GB of swap is recommended. Creating a large swap space partition can be especially helpful if you plan to upgrade your RAM at a later time.

Tip

If your partitioning scheme requires a swap partition that is larger than 2GB, you should create an additional swap partition. For example, if you have 4GB of RAM, you may want to create two 2GB swap partitions.

- A PReP boot partition on the first partition of the hard drive - the PReP boot partition contains the operating system kernel or the yaboot boot loader (which allows pSeries systems to boot Red Hat Enterprise Linux). On pSeries, unless you plan to boot from a floppy or network source, you must have a PReP boot partition to boot Red Hat Enterprise Linux. The PReP boot partition should be between 4-8MB for pSeries systems, not to exceed 10MB.
- A /boot/ partition (100 MB) - the partition mounted on /boot/ contains the operating system kernel (which allows your system to boot Red Hat Enterprise Linux), along with files used during the bootstrap process. Due to the limitations of most PC firmware, creating a small partition to hold these is a good idea. For most users, a 100MB boot partition is sufficient.

Caution

If you have a RAID card, be aware that some firmware do not support booting from the RAID card. In cases such as these, the /boot/ partition must be created on a partition outside of the RAID array, such as on a separate hard drive.

- A root partition (350MB - 5.0GB) - this is where "/" (the root directory) is located. In this setup, all files (except those stored in /boot) are on the root partition. A 350MB partition allows you to install a minimal installation, while a 5.0GB root partition lets you perform a full installation, choosing all package groups.
- A /var/ partition (3.0GB or larger) - the /var/ partition is where variable data files are written. This includes spool directories and files, administrative and logging data, and transient and temporary files. Updates that are applied to Red Hat Enterprise Linux are written to /var/spool/update/ and require approximately 1.0GB of disk space available when updating your system.

After completing the disk partitioning step, the next installation program screen will allow you to set your Network Configuration for eth0. Disable the eth0 network device by deselecting "Activate on boot". Click on OK and then configure eth1. The eth1 configuration information should be remembered from earlier, for example, blade1's configuration information is: IP Address of 192.168.70.100, and a Netmask of 255.255.255.0. Click on OK. Next, some Miscellaneous Network Settings are shown. Again, these should already be filled in. For example, blade1's miscellaneous settings would be: Gateway of 192.168.70.1 and the Primary DNS of the network. Click OK to continue. The network configuration can be changed after the installation by running the "redhat-config-network" command.

The last piece of networking information is shown on the Hostname Configuration screen. Select to manually provide a hostname, for blade1 this would be blade1.bclab.com. Click OK to continue.



Next is the Firewall Configuration screen. We'll select "No firewall" since the servers are usually connected to a router providing firewall services. Click on OK. The firewall configuration can be changed after the installation by running the "redhat-config-securitylevel" command.

Then the Language Support Selection screen is shown. Choose any additional languages to install and support on your system. Click OK when finished.

The next screen will be the Timezone Configuration screen. Select the appropriate timezone for the server. The timezone configuration can be changed after the installation by running the "redhat-config-date" command. Click OK to continue.

Next, you'll be given the option to set the root password. The root password can be changed after the installation by running the "redhat-config-rootpassword" command. Click OK when done.

The last installation program screen is the Package Defaults screen, which allows you to select the software packages to be installed. You can accept the recommended package selection or customize the software selection. To accept the recommended selection, simply click on OK. To customize, select Customize software selection, and click on OK. When finished customizing, click on OK. Now, the installation of RHEL will begin. The installation takes a variable amount of time depending on the quantity and size of the packages selected.

While the installation is proceeding, the Package Installation screen will show the current package being installed and a running total of the installation packages and time estimate.

When the installation completes, you'll be shown the Complete screen. Prior to rebooting you'll need to set the boot sequence now that installation is complete. Start the MM Web interface, as described in *Section I, Using the Management Module's Web interface*. Under "Blade Tasks", select "Configuration", then "Boot Sequence". Select the server to modify. Insert "Hard drive 0" first in the startup sequence for the server(s) just installed and click on "Save".

Now, click on OK on the Complete screen. The system will reboot using the new installation.

When yaboot presents the "boot:" prompt, either press Enter or let it timeout. In either case it will boot the default boot entry. After one or more screens of messages scroll by, a "login:" prompt will be shown. Login as the root user and supply the root password provided during the configuration. The installation is now complete.

IV. Auto-installation over network via OpenFirmware

This install procedure is best used when frequently installing many systems with a similar configuration or categories of configurations (i.e. utilizing the kickstart configuration file). In this procedure, we'll pass kernel options to the install kernel by using the OpenFirmware boot command. This install procedure requires:

- Setup of the install server
 1. DHCP configuration
 2. TFTP configuration
 3. NFS configuration
 4. Building the installation tree



5. Create the kickstart configuration file
 - Install from OpenFirmware

Setup of the install server

Setup DHCP, TFTP, NFS and the install tree as was done in *Section III, Setup of the install server*.

Create the kickstart configuration file

As a final piece of install server setup, create a kickstart configuration file which will contain the predefined “answers” to the configuration options and be used in lieu of manually stepping through the Installation Program screens. These are suggested alternatives for generating the kickstart file:

1. Use the sample.ks file found in the RH-DOCS directory of the Red Hat Enterprise Linux Documentation CD
2. Reuse a configuration from an existing RHEL system as a starting point
3. Generate a configuration file using the Kickstart Configurator application
4. A sample JS20 kickstart configuration file is available for download from the same location that you downloaded this document from.

Check the sample.ks file in the RH-DOCS directory of the RHEL Documentation CD. For a detailed description of all the kickstart options, see Chapter 9 of the *Red Hat Enterprise Linux System Administration Guide*.

The RHEL installation program creates a sample kickstart file based on the options that were selected during installation in `/root/anaconda-ks.cfg`. You should be able to find this file on any RHEL system.

To be guided through the generation of a kickstart configuration file, run:

```
# redhat-config-kickstart
```

This provides a GUI that intelligently creates a kickstart configuration file.

Once a configuration file is available, it should be saved on the install server with a filename of the form `<ip-addr>-kickstart`. For example, for blade1 our kickstart file would be saved in:

```
/install-tree/kscfg/192.168.70.100-kickstart
```

Based on our installation tree setup, a line that is required in our kickstart configuration file is a pointer to the installation tree, such as:

```
nfs --server 192.168.70.50 --dir /install-tree/rhelas3/QU3
```

Install from OpenFirmware



When installing from OpenFirmware it is not necessary to set the boot sequence in the MM Web interface.

To reboot the server, and thereby kick off the automatic installation process, under “Blade Tasks” on the left hand side of the MM Web interface, select “Power/Restart”. Check the box for the blade you wish to start installing and power it on by selecting “Power On Blade” or “Restart Blade”.

Start an SOL session as described in *Section 1, Using a Serial Over LAN session*, in order to view the installation. On the SOL console, you should see the progress codes scrolling past while the JS20 initializes. When starting the install via OpenFirmware, the SOL session should not be interrupted.

When you see the **D5BB progress code**, it will stay there for 5 seconds and display a countdown. While this is counting down from 5 to 1, press 8 on the keyboard. This will take you into OpenFirmware, the “O >” prompt. At the OpenFirmware prompt, run the “ls” command, the output should contain a /pci section that looks like:

```
O > ls
.....
000000d34278: /pci@8000000f8000000
000000d37d28: /pci@0
000000d4b408: /ethernet@1
000000d5a788: /ethernet@1,1
.....
ok
```

We need to issue the following “boot” command:

```
O > boot /pci@8000000f8000000/pci@0/ethernet@1,1:bootp,192.168.70.50,,192.168.70.100,0.0.0.0
ks=nfs:192.168.70.50:/install-tree/kscfg/192.168.70.100-kickstart ksdevice=eth1
```

Where each of the fields is defined as:

- full path name of eth1 device = /pci@8000000f8000000/pci@0/ethernet@1,1
 - issue the “devalias” command to see if there is an alias created for the 2nd Ethernet adapter. If there isn’t already an alias, you can create one by using the devalias command. For example, “devalias eth1 /pci@8000000f8000000/pci@0/ethernet@1,1” will create an alias of eth1 for the 2nd Ethernet adapter.
- arguments to Ethernet component:
 - protocol = bootp
 - server IP address = 192.168.70.50
 - client IP address = 192.168.70.100
 - gateway IP address = 196.168.70.1
- arguments passed to the install kernel are:
 - location of kickstart configuration file = “ks=nfs:192.168.70.50:/install-tree/kscfg/192.168.70.100-kickstart”
 - directive for JS20 to install over eth1 = ksdevice=eth1

The OpenFirmware prompt appears to provide **no line wrap feature**, so just continue typing when you get to the end of line. You’ll see the whole command there if you need to retrieve it. For boot command details, see <http://www.firmworks.com/QuickRef.html>.



The installation process is visible via the SOL session. When the installation starts, there will be screens shown about acquiring a DHCP address, formatting the hard drive and installing packages. Of course, these screens are driven by the content of the kickstart configuration file. Once the installation is complete, a "login:" prompt will be displayed on the SOL session. The password will be the password specified in the kickstart configuration file. At this point, the installation is complete.

V. Auto-installation over network using CSM

This section describes how to populate the blade servers using Cluster Systems Management (CSM) for Linux V1.4.0.3. CSM for Linux V1.4.0.3 is the first release of CSM to officially support RHEL AS 3 Update 3. Using CSM for installs is most useful when you want the blade servers to operate as a cluster of nodes. CSM is an administration tool that will provide an install server capability as well as features that simplify the ongoing administration tasks of the cluster. For detailed descriptions of those administration features, see the *CSM for Linux V1.4 Administration Guide*. This section will address the OS installation capabilities of CSM.

In the rest of this section we'll assume the CSM management server (ie. the install server) is a separate system (not a blade). This is an outline of the steps necessary to perform this type of installation:

- CSM management server (MS) setup
- Define the nodes to the MS
- CSM Red Hat (kickstart) installation setup
- Install the OS to the nodes
- Monitoring with CSM - optional

CSM management server (MS) setup

The CSM MS installation is an RPM-based installation made up of these steps:

1. Setup /etc/hosts
2. Download the prerequisite RPMS, in this case Autoupdate and OpenCIMOM
3. Install the csm.core RPM
4. Invoke the CSM command "installms", which installs the remaining RPMS
5. Accept the product license
6. Verify the CSM MS installation

Prior to installing the RPMS, we'll setup /etc/hosts for hostname resolution, here's an example:

```
# hosts
```

```
192.168.70.50      ms.bclab.com      ms
192.168.70.125    mm.bclab.com      mm
192.168.70.126    mm-int.bclab.com  mm-int
192.168.70.130    esm1.bclab.com    esm1
192.168.70.131    esm2.bclab.com    esm2
```



```
192.168.70.100    blade1.bclab.com    blade1
192.168.70.101    blade2.bclab.com    blade2
```

```
# end of hosts
```

You'll also need to obtain the Autoupdate RPM and download it to a temporary directory, for example: /tmp/csmreqs. The Autoupdate RPM can be downloaded from:

<ftp://ftp.istrata.com/pub/autoupdate/>.

Download the OpenCIMOM RPM in a similar manner. The OpenCIMOM RPM can be found at: <ftp://ftp.software.ibm.com/aix/freeSoftware/aixtoolbox/RPMS/ppc/openCIMOM/openCIMOM-0.8-1.noarch.rpm>

Proceed with **one of the next two** sections, depending on whether the CSM software package will be downloaded from the Web or CSM product CDs will be used.

CSM software package from Website

The first RPM to install is the csm.core RPM. If you're obtaining the CSM package as a tar file from the Website (<http://techsupport.services.ibm.com/server/cluster/fixes>), download it to the /tmp/csm directory on the install server and uncompress using:

```
# cd /tmp/csm
# tar -xvzf csm-linux-1.4.0.3.ppc64.tar.gz
```

Install the csm.core RPM by running:

```
# rpm -ivh /tmp/csm/csm.core-*
```

Installing the csm.core RPM will alter the \$PATH and \$MANPATH environment variables. To activate these environment variables, exit the current shell and start a new shell (logout and log back in).

Now we need to install the other RPMs from the CSM package on the install server using the installms command from the csm.core RPM:

```
# installms -p /tmp/csm:/tmp/csmreqs
```

Then insert the RHEL CDs when prompted. The CDs required, when installing CSM on a RHEL system are:

- RHELAS3 PPC (Update 1) disk1
- RHELAS3 PPC disk2
- RHELAS3 PPC disk3
- RHELAS3 PPC (Update 1) Extras disk1
- RHELAS3 PPC (Update 2) disk4

When prompted about potentially conflicting TFTP packages, select option 1 to remove the existing TFTP packages and install CSM's TFTP package.



The installation of CSM on the MS should complete with this message "Installation of CSM has successfully completed!".

CSM software package from CDs

The first RPM to install is the csm.core RPM. If you've obtained the CSM package as a CD, install the csm.core RPM by running:

```
# mount /dev/cdrom /mnt/cdrom
# rpm -ivh /mnt/cdrom/csm.core.*
```

Installing the csm.core RPM will alter the \$PATH and \$MANPATH environment variables. To activate these environment variables, exit the current shell and start a new shell (logout and log back in).

Now we need to install the other RPMs from the CSM package on the install server using the installms command from the csm.core RPM:

```
# installms -p /mnt/cdrom:/tmp/csmreqs
```

Then insert the RHEL CDs when prompted. The CDs required, when installing CSM on a RHEL system are:

- RHELAS3 PPC (Update 1) disk1
- RHELAS3 PPC disk2
- RHELAS3 PPC disk3
- RHELAS3 PPC (Update 1) Extras disk1
- RHELAS3 PPC (Update 2) disk4

When prompted about potentially conflicting TFTP packages, select option 1 to remove the existing TFTP packages and install CSM's TFTP package.

The installation of CSM on the MS should complete with this message "Installation of CSM has successfully completed!".

Accept the product license

Next, we need to either: accept the full CSM license from the CSM CD or use the 60 day try-and-buy CSM license. To accept the license from the product CD, insert the CSM CD into the CD-ROM drive and run these commands:

```
# mount /dev/cdrom /mnt/cdrom
# csmconfig -L /mnt/cdrom/csmlum.full
```

Verify the full license is in place by running:

```
# csmconfig
```

The output should look like this:

```
AddUnrecognizedNodes = 0 (no)
```



```
ClusterSNum = 654BN5011
ClusterTM = 9078-160
DeviceStatusFrequency = 12
DeviceStatusSensitivity = 8
ExpDate =
HAMode = 0
HeartbeatFrequency = 12
HeartbeatSensitivity = 8
LicenseProductVersion = 1.4
MaxNumNodesInDomain = -1 (unlimited)
PowerStatusMode = 0 (Mixed)
RegSyncDelay = 1
RemoteCopyCmd = /usr/bin/scp
RemoteShell = /usr/bin/ssh
SetupKRB5 = 0
SetupRemoteShell = 1 (yes)
```

To accept the 60 day try-and-by license, run:

```
# csmconfig -L
```

Verify the try-and-buy license is in place by running:

```
# csmconfig
```

The output should look like this:

```
AddUnrecognizedNodes = 0(no)
ClusterSNum = 654BN5011
ClusterTM = 9078-160
ExpDate = Tue Oct 19 18:59:59 2004
HeartbeatFrequency = 12
HeartbeatSensitivity = 8
LicenseProductVersion = 1.4
MaxNumNodesInDomain = -1 (unlimited)
PowerstatusMode = 0(Mixed)
RegSyncDelay = 1
RemoteCopyCmd = /usr/bin/scp
RemoteShell = /usr/bin/ssh
SetupKRB5 = 0
SetupRemoteShell = 1(yes)
```

Note the highlighted fields, those fields indicate the appropriate license has been accepted and is in place. Prior to accepting one of the licenses, the setting of "MaxNumNodesInDomain = 0" would prohibit defining any nodes to the MS. After accepting one of the licenses, "MaxNumNodesInDomain = -1 (unlimited)" is the new setting.

To verify the installation of the CSM MS was successful, run:

```
# probemgr -p ibm.csm.ms -l 0
```

The last line of the probemgr output should be "Probe ibm.csm.ms was run successfully."



Define the Nodes to the MS

In order to install the OS to a blade server, we must define the blade server as a “node” to the CSM MS. This will be accomplished in the following steps:

1. gather attributes from the HW control point using the `Ishwinfo` command
2. make minor alterations to the `Ishwinfo` output file
3. transform the `Ishwinfo` output file to a node definition file format (ie. `definenode .. -s ..`)
4. define the nodes, using the node definition file as input (ie. `definenode .. -f ..`)

The BladeCenter's MM is considered the hardware control point of the cluster. In CSM terminology, that means the MM's IP address (or hostname) is to be used as the `HWControlPoint` attribute. We'll gather as many node attributes, automatically, using the HW control capabilities of the MM. Issue this command:

```
# Ishwinfo -p blade -c mm -o /tmp/csm/hwinfo1
```

The resulting `/tmp/csm/hwinfo1` file should look like this:

```
# Hostname::PowerMethod::HWControlPoint::HWControlNodeld::LParID::HWType::HWModel::HWSerialNum  
no_hostname::blade::mm.bclab.com::blade1:::8842::41X::ZJ1W6745E17J  
no_hostname::blade::mm.bclab.com::blade2:::8842::41X::ZJ1W6745E17G
```

Note, the first line of the `hwinfo1` file provides field headings for the information contained in the subsequent lines. Each field is delineated by two colons, “:”.

Next, we'll edit the `hwinfo1` file to supply actual hostnames for blade servers in our environment. Edit the `hwinfo1` file and substitute the actual hostname for the “`no_hostname`” entry in each row. The resulting file should resemble:

```
# Hostname::PowerMethod::HWControlPoint::HWControlNodeld::LParID::HWType::HWModel::HWSerialNum  
blade1::blade::mm.bclab.com::blade1:::8842::41X::ZJ1W6745E17J  
blade2::blade::mm.bclab.com::blade2:::8842::41X::ZJ1W6745E17G
```

Now we'll use CSM's `definenode` command to create a node definition file based on `hwinfo1`. At the same time, we'll prime the resulting node definition file with attribute/value pairs that need to be specified for all JS20 nodes. Note, the attribute/value pairs are at the end of the `definenode` command. Run this `definenode` command:

```
# definenode -M hwinfo1 -C mm.bclab.com::1:2 -s ConsoleMethod=blade \  
InstallAdapterName=eth1 InstallAdapterDuplex=full InstallAdapterSpeed=1000 \  
InstallAdapterType=ent InstallDiskType=ide ConsoleSerialDevice= > nodedef1
```

Note, the “`InstallDistributionName`” and “`InstallDistributionVersion=`” `attr=value` pairs may be necessary if the CSM MS is not running the same distribution or version as the installation target. In our scenario, the CSM MS is a RHEL AS 3 Update 3 system, so it is not necessary to specify those attributes.

The resulting `nodedef1` file should look like this (with important JS20 attributes in boldface):

```
blade1.bclab.com :  
  ConsoleMethod=blade  
  ConsolePortNum=1
```



```
ConsoleSerialDevice=  
ConsoleServerName=mm.bclab.com  
HWControlNodeId=blade1  
HWControlPoint=mm.bclab.com  
HWModel=41X  
HWSerialNum=ZJ1W6745E17J  
HWType=8842  
InstallAdapterDuplex=full  
InstallAdapterName=eth1  
InstallAdapterSpeed=1000  
InstallAdapterType=ent  
InstallCSMVersion=1.4.0  
InstallDiskType=ide  
InstallDistributionName=RedHatEL-AS  
InstallDistributionVersion=3  
InstallOSName=Linux  
InstallPkgArchitecture=ppc64  
ManagementServer=ms.bclab.com  
PowerMethod=blade  
blade2.bclab.com :  
ConsoleMethod=blade  
ConsolePortNum=2  
ConsoleSerialDevice=  
ConsoleServerName=mm.bclab.com  
HWControlNodeId=blade2  
HWControlPoint=mm.bclab.com  
HWModel=41X  
HWSerialNum=ZJ1W6745E17G  
HWType=8842  
InstallAdapterDuplex=full  
InstallAdapterName=eth1  
InstallAdapterSpeed=1000  
InstallAdapterType=ent  
InstallCSMVersion=1.4.0  
InstallDiskType=ide  
InstallDistributionName=RedHatEL-AS  
InstallDistributionVersion=3  
InstallOSName=Linux  
InstallPkgArchitecture=ppc64  
ManagementServer=ms.bclab.com  
PowerMethod=blade
```

Next, we'll feed the `nodedef1` file to CSM's `definenode` command by running this command:

```
# definenode -f nodedef1
```

Upon completion of the `definenode` command, we can run:

```
# lsnode -l
```

... to get a list of all the nodes and their attributes, should result in this output:



```
Hostname = blade1.bclab.com
AllowManageRequest = 0 (no)
CSMVersion =
ChangedAttributes = {}
ConfigChanged = 0 (no)
ConsoleMethod = blade
ConsolePortNum = 1
ConsoleSerialDevice =
ConsoleSerialSpeed = 9600
ConsoleServerName = mm.bclab.com
ConsoleServerNumber =
HWControlNodeID = blade1
HWControlPoint = mm.bclab.com
HWModel = 41X
HWSerialNum = ZJ1W6745E17J
HWType = 8842
InstallAdapterDuplex = full
InstallAdapterGateway =
InstallAdapterMacaddr =
InstallAdapterName = eth1
InstallAdapterNetmask =
InstallAdapterSpeed = 1000
InstallAdapterType = ent
InstallCSMVersion = 1.4.0
InstallDisk =
InstallDiskType = ide
InstallDistributionName = RedHatEL-AS
InstallDistributionVersion = 3
InstallKernelVersion =
InstallMethod =
InstallOSName = Linux
InstallPkgArchitecture = ppc64
InstallServer =
InstallStatus = PreManaged
InstallTemplate =
LParID =
LastCFMUpdateTime =
ManagementServer = ms.bclab.com
Mode = PreManaged
Name = blade1.bclab.com
NodeNameList = {ms}
PowerMethod = blade
PowerStatus = 1 (on)
Status = 127 (unknown)
UpdatenodeFailed = 0 (false)
UserComment =
-----
Hostname = blade2.bclab.com
AllowManageRequest = 0 (no)
CSMVersion =
ChangedAttributes = {}
ConfigChanged = 0 (no)
```




```
ConsoleMethod = blade
ConsolePortNum = 2
ConsoleSerialDevice =
ConsoleSerialSpeed = 9600
ConsoleServerName = mm.bclab.com
ConsoleServerNumber =
HWControlNodeID = blade2
HWControlPoint = mm.bclab.com
HWModel = 41X
HWSerialNum = ZJ1W6745E17G
HWType = 8842
InstallAdapterDuplex = full
InstallAdapterGateway =
InstallAdapterMacaddr =
InstallAdapterName = eth1
InstallAdapterNetmask =
InstallAdapterSpeed = 1000
InstallAdapterType = ent
InstallCSMVersion = 1.4.0
InstallDisk =
InstallDiskType = ide
InstallDistributionName = RedHatEL-AS
InstallDistributionVersion = 3
InstallKernelVersion =
InstallMethod =
InstallOSName = Linux
InstallPkgArchitecture = ppc64
InstallServer =
InstallStatus = PreManaged
InstallTemplate =
LParID =
LastCFMUpdateTime =
ManagementServer = ms.bclab.com
Mode = PreManaged
Name = blade2.bclab.com
NodeNameList = {ms}
PowerMethod = blade
PowerStatus = 1 (on)
Status = 127 (unknown)
UpdatenodeFailed = 0 (false)
UserComment =
```

Note, if you need to change attributes for a node for any reason, you'll need to use the CSM `chnode` command. For a list of all the possible node attributes, issue the command "man `nodeattributes`". Note that "nodeattributes" is not a CSM command, there is just a man page defined to describe the node attributes.



CSM Red Hat (kickstart) installation setup

Now we need to supply the RHEL CDs to CSM so it can build the install tree and the kickstart configuration file.

Unlike other install procedures, where we've had to alter the boot sequence before and after installation, CSM will fix up the boot sequence after installation. We need to have Hard drive 0 as the first entry in the boot sequence to start the process. Verify this by starting the MM Web interface as described in *Section I, Using the Management Module's Web interface*. Under "Blade Tasks", select "Configuration", then "Boot Sequence". Select the server to be installed. Verify that "Hard drive 0" is first in the boot sequence. If it isn't, select "Hard drive 0" as the first entry in the sequence and click on "Save".

Note, at this point you can verify remote power control of the nodes. You should be able to enter:

```
# rpower -a query
```

and see results similar to:

```
blade1.bclab.com on  
blade2.bclab.com off
```

If results do not look like that, run:

```
# probemgr -p ibm.csm.HWCtrl
```

CSM maintains templates for the different automated installs it performs. If you want to alter the kickstart configuration file template for the upcoming install, the template is located in `/opt/csm/install/kscfg.tmpl.RedHatEL-AS3-ppc64`. For further details on creating/modifying/customizing a kickstart control file see *Section II, Red Hat Enterprise Linux System Administration Guide*.

Next, we need to supply the RHEL AS 3 Update 3 CDs to CSM, run the following command:

```
# copycds InstallDistributionName=RedHatEL-AS InstallDistributionVersion=3  
InstallServiceLevel=QU3
```

and insert the next RHEL AS 3 Update 3 CD when prompted. As CSM copies the contents of the CD, the % complete should be displayed in the lower left of the window.

CSM maintains the content of these CDs in `/csminstall/Linux/RedHatEL-AS/3/ppc64`.

Prior to issuing the CSM commands of `csmsetupks` and `installnode`, run:

```
# monitorinstall
```

to verify the state of each node. The output should look like this:

Node	Mode	Status
blade1.bclab.com	PreManaged	Not Installed
blade2.bclab.com	PreManaged	Not Installed



Next, we need to run:

```
# csmsetupks -x -P
```

to create a kickstart configuration file for each of the nodes in the cluster. The “-P” option indicates to create the control file for each node with Mode=PreManaged. The “-x” option indicates the CD images have already been copied. The output of the csmsetupks command can be found in /var/log/csm/csmsetupks.log. Most of the CSM commands provide the following options to identify the target(s) of the command:

- -P: for all nodes where Mode=Premanaged
- -a: all nodes
- -n: followed by a node or comma separated list of nodes (eg. -n blade1)
Also, see the noderange man page for a shorthand way to specify many nodes.
- -N: followed by a nodegroup or comma separate list of nodegroups (eg. -N BladeNodes)

CSM maintains the content of the node’s actual kickstart configuration files in /csminstall/csm/1.4.0/kickstart.RedHatEL-AS3.

Install the OS to the nodes

The next step is to install RHEL AS 3 Update 3 and the CSM client code on the nodes, to do this run:

```
# installnode -P
```

This is an asynchronous command. To monitor the installation status, run:

```
# monitorinstall
```

typical output might be:

Node	Mode	Status
blade1.bclab.com	Installing	Rebooting and Installing Node.
blade2.bclab.com	Managed	makenode complete

To continuously monitor the installation, run:

```
# watch monitorinstall
```

When finished installing the nodes, the monitorinstall results will be:

Node	Mode	Status
blade1.bclab.com	Managed	Installed
blade2.bclab.com	Managed	Installed

When the installation is complete, you can login to the nodes using CSM’s default password of “cluster”.



To verify that Resource Monitoring and Control (RMC) is working on the nodes, run:

```
# lsnode -H -l
```

The output should look similar to:

```
Name = blade1
ActiveMgtScopes = 1
Architecture = ppc64
DistributionName = RedHatEL-AS
DistributionVersion = 3
KernelVersion = 2.4.21-147-pseries64
LoadAverage = {0,0,0.01}
NodeNameList = {blade1.bclab.com}
NumOnlineProcessors = 2
NumProcessors = 2
NumUsers = 1
OSName = Linux
PctRealMemFree = 29
PctTotalPgSpFree = 100
PctTotalPgSpUsed = 0
PctTotalTimeIdle = 96.2245
PctTotalTimeKernel = 1.06353
PctTotalTimeUser = 2.67518
PctTotalTimeWait = 0.0367596
RealMemSize = 477413376
TotalPgSpFree = 155734
TotalPgSpSize = 155734
UpTime = 2137
VMPgInRate = 0
VMPgOutRate = 8
VMPgSpInRate = 0
VMPgSpOutRate = 0
-----
Name = blade2
ActiveMgtScopes = 1
Architecture = ppc64
DistributionName = RedHatEL-AS
DistributionVersion = 3
KernelVersion = 2.4.21-147-pseries64
LoadAverage = {0,0,0}
NodeNameList = {blade2.bclab.com.com}
NumOnlineProcessors = 2
NumProcessors = 2
NumUsers = 1
OSName = Linux
PctRealMemFree = 29
PctTotalPgSpFree = 100
PctTotalPgSpUsed = 0
PctTotalTimeIdle = 96.4853
PctTotalTimeKernel = 1.03231
```



```
PctTotalTimeUser = 2.44684
PctTotalTimeWait = 0.0355847
RealMemSize = 477409280
TotalPgSpFree = 155734
TotalPgSpSize = 155734
UpTime = 2341
VMPgInRate = 0
VMPgOutRate = 51
VMPgSpInRate = 0
VMPgSpOutRate = 0
```

You can also verify the Distributed Shell (dsh) command is working, run:

```
# dsh -as date
```

The output should look similar to:

```
blade1.bclab.com: Wed Aug 18 22:06:02 CDT 2004
blade2.bclab.com: Wed Aug 18 22:08:17 CDT 2004
```

Monitoring with CSM - optional

CSM contains some easy-to-use monitoring commands. In CSM, monitoring of an event begins when a "condition" and "response" pair is activated. On the MS, there are some predefined conditions immediately available, view them by running the command:

```
# lscondition
```

Similarly, the predefined responses can be listed by running the command:

```
# lsresponse
```

The currently activated condition and response pairs can be listed by running:

```
# lscondresp
```

The listing below shows the automatically activated condition and response pairs:

Displaying condition with response information:

Condition	Response	Node	State
"NodeFullInstallComplete"	"RunCFMToNode"	"ms"	"Active"
"NodeFullInstallComplete"	"GatherKRB5keytabs"	"ms"	"Active"
"pSeriesNodeFullInstallComplete"	"removeArpEntries"	"ms"	"Active"
"NodeFullInstallComplete"	"GatherSSHHostKeys"	"ms"	"Active"
"UpdatenodeFailedStatusChange"	"UpdatenodeFailedStatusResponse"	"ms"	"Active"
"NodeChanged"	"consoleUpdateResponse"	"ms"	"Active"

We'll monitor the event of a blade server changing its power state (to either on or off) and log the event. To this we'll use the predefined condition `NodePowerStatus` and the predefined condition `LogCSMEventsAnyTime`. To activate this condition and response pair, run this command:



```
# startcondresp NodePowerStatus LogCSMEventsAnyTime
```

The output of:

```
# lscondresp
```

should now contain this additional event pairing:

Condition	Response	Node	State
"NodePowerStatus"	"LogCSMEventsAnyTime"	"ms"	"Active"

The most important attribute of a response is its action script. To see the attributes that define the LogCSMEventsAnyTime response, run:

```
# lsresponse LogCSMEventsAnyTime
```

This is the output that will be shown:

Displaying response information:

```
ResponseName = "LogCSMEventsAnyTime"
Node         = "ms"
Action       = "logEvent"
DaysOfWeek   = 1-7
TimeOfDay    = 0000-2400
ActionScript = "/usr/sbin/rsct/bin/logevent /var/log/csm/systemEvents"
ReturnCode   = 0
CheckReturnCode = "n"
EventType    = "b"
StandardOut  = "n"
EnvironmentVars = ""
UndefRes     = "n"
```

The LogCSMEventsAnyTime response will add a log entry to the file `/var/log/csm/systemEvents` whenever one of the blades is powered off and on.

Users can make their own conditions and responses as well. For details of the CSM monitoring capabilities, see the *CSM for Linux V1.4 Administration Guide*.

VI. Red Hat Network

The Red Hat Network (RHN) is an Internet solution for managing RHEL systems. All Security Alerts, Bug Fix Alerts, and Enhancement Alerts (collectively known as Errata Alerts) can be downloaded directly from Red Hat. You can even have updates scheduled to be delivered directly to your system as soon as they are released.

RHN components

The main components of RHN are as follows:



- the Red Hat Update Agent
- the Red Hat Network Web site
- Red Hat Network Daemon

The Red Hat Update Agent (up2date) provides your initial connection to Red Hat Network. Once registered, it enables channel subscription, package installs, and management of System Profiles. Red Hat Enterprise Linux AS 3 and newer systems will even use the Red Hat Update Agent to be registered to RHN.

Although both the RHN Web site and the Red Hat Update Agent allow you to view Errata Alerts from the Red Hat Errata list and apply updates, the Website offers a much more granular view of your account and provides added functionality required for management and provisioning of systems. For enterprise deployments of Red Hat Network, the Web site is a must.

The Red Hat Network Daemon (rhnsd) runs in the background as a service and probes the Red Hat Network for notifications and updates at set time intervals. This daemon is necessary if you want to schedule updates or other actions through the Web site.

RHN service levels

RHN lets the user choose among three levels of service: Update, Management and Provisioning. For details on each level of service, see: <http://www.redhat.com/software/rhn/table/> In the remainder of this section, we'll focus on the Update service level.

RHN Update service overview

The RHN Update service level is targetted for users with one, or a small number of RHEL systems to maintain. The RHN Update service is thoroughly documented in the RHN Update Reference Guide (<https://rhn.redhat.com/help/basic/index.html>). A brief overview follows.

First you need to register your system, run:

```
# up2date -- register
```

Upon registering, you'll be asked if you already have a user or organizational account. If so, it's recommended that you add systems to an existing account, instead of creating new accounts for each system. Next, you'll register your systems hardware and package profiles and send them to Red Hat.

After finishing registration, the next step is to select a channel from which to receive updates. Once the channel is selected, a connection RHN is made and your customized list (ie. not the updates themselves, just a list) of updates is retrieved. Next, you'll be shown a variety of screens that let you select previously uninstalled packages, or packages that have been updated. You can also view information about each package, its Bug Fix Advisory and select the packages you wish to update. After making the selections, the Red Hat Update Agent tests the package dependencies and prompts the user for any recommended changes. After testing the updates are downloaded. By default, the Red Hat Update Agent will also install packages it has downloaded.



There are many other features to the RHN Update service. The Red Hat Update Agent configuration tool can be used to customize the update process. The Red Hat Network Alert Notification tool is a notifier that appears on the GNOME panel.

VII. Troubleshooting

Serial Over LAN problems

If Serial Over LAN sessions aren't showing you the expected output, here are some things to try.

1. Disable and re-enable SOL on the JS20. To do this, under "Blade Tasks", select "Serial Over LAN". Check the box next to the appropriate JS20. Click on the "Disable Serial Over LAN" link at the bottom. After the page reloads, the SOL column of the table should indicate "Disabled". Then check the JS20's box again and click on the "Enable Serial Over LAN" link at the bottom. After the page reloads, the SOL column of the table should indicate "Enabled".
2. Restart the SP on the JS20. To do this, under "Blade Tasks", select "Power/Restart". Check the box next to the appropriate JS20. Click on the "Restart the Blade System Management Processor" link at the bottom, and confirm by pressing OK. After a short period of time, retry the SOL session.
3. As a last resort, remove the JS20 from the BladeCenter unit, then reinsert the JS20 into the BladeCenter unit. Make sure to reinsert the JS20 into the same slot from which it was removed. Follow the instructions for this procedure as stated in the BladeCenter JS20 Installation and User's Guide.

DHCP problems

Make sure there is only one DHCP server running on the network. The install server (or CSM MS) should have DHCPDISCOVER, DHCPOFFER, DHCPREQUEST, DHCPACK log entries for each install target in /var/log/messages.



TFTP problems

If you encounter problems TFTPing the install kernel from the install server, it may be the TFTP daemon on the install server is not capable of handling a number of parallel sessions. In this case, replace the TFTP package on the server with the atftp RPM.

Review the notes in *Section III, DHCP Configuration* and *TFTP Configuration* about specifying the accurate filename in `dhcpd.conf` and the proper file ownership in `/tftpboot`.

CSM problems

If the installation of CSM on the management server does not complete successfully, run:

```
# probemgr -p ibm.csm.ms -l 0
```

There are a number of probes available in CSM. For details, see the CSM for Linux V1.4 Administration Guide.

CSM maintains its log files in `/var/log/csm`. These files can provide per-command-instance information (eg. `installms.log.1`, `installms.log.2`, etc...).

To monitor discussion of CSM topics, subscribe to one of the CSM mailing lists.

- the external list: <http://www-124.ibm.com/developerworks/oss/mailman/listinfo/csm>
- the internal list: <http://lxsvr4.ppd.pok.ibm.com/cgi-bin/lwgate/CSMDEV/>



VIII. References

1. Red Hat Enterprise Linux Installation Guide for the IBM eServer iSeries and IBM eServer pSeries Architectures,
<http://www.redhat.com/docs/manuals/enterprise/RHEL-3-Manual/ppc-multi-install-guide/>
2. Red Hat Enterprise Linux System Administration Guide,
<http://www.redhat.com/docs/manuals/enterprise/RHEL-3-Manual/sysadmin-guide/>
3. BladeCenter JS20 Type 8842 Installation and User's Guide,
<http://www-307.ibm.com/pc/support/site.wss/document.do?Indocid=MIGR-54526>
4. IBM eServer BladeCenter - Management Module User's Guide,
<http://www-307.ibm.com/pc/support/site.wss/document.do?Indocid=MIGR-45153>
5. IBM eServer BladeCenter – Planning and Installation Guide,
<http://www-307.ibm.com/pc/support/site.wss/document.do?Indocid=MIGR-53670>
6. Serial over LAN Setup Guide – IBM eServer BladeCenter and BladeCenter T,
<http://www-307.ibm.com/pc/support/site.wss/document.do?Indocid=MIGR-54666>
7. CSM for Linux 1.4 Administration Guide,
<http://publib.boulder.ibm.com/clresctr/docs/csm/linux/200408/am7al103/am7al10302.html>
8. CSM for Linux 1.4 Planning and Installation Guide,
<http://publib.boulder.ibm.com/clresctr/docs/csm/linux/200408/am7il103/am7il10302.html>
9. CSM for Linux 1.4 Command and Technical Reference,
<http://publib.boulder.ibm.com/clresctr/docs/csm/linux/200408/am7cl102/am7cl10202.html>
10. RHN Update Reference Guide, <https://rhn.redhat.com/help/basic/index.html>



© Copyright IBM Corporation 2004

IBM Corporation
Marketing Communications
Systems and Technology Group
Route 100
Somers, New York 10589

Produced in the United States of America
October 2004
All Rights Reserved

This document was developed for products and/or services offered in the United States. IBM may not offer the products, features, or services discussed in this document in other countries.

The information may be subject to change without notice. Consult your local IBM business contact for information on the products, features and services available in your area.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

IBM, the IBM logo, @server, eServer, BladeCenter, iSeries, pSeries and ThinkPad are trademarks or registered trademarks of International Business Machines Corporation in the United States or other countries or both. See <http://www.ibm.com/legal/copytrade.shtml>.

Linux is a registered trademark of Linus Torvalds in the United States, other countries or both.

Windows is a registered trademark of the Microsoft Corporation.

Other company, product, and service names may be trademarks or service marks of others.

This equipment is subject to FCC rules. It will comply with the appropriate FCC rules before final delivery to the buyer.

IBM hardware products are manufactured from new parts, or new and used parts. Regardless, our warranty terms apply.

Information concerning non-IBM products was obtained from the suppliers of these products. Questions on the capabilities of the non-IBM products should be addressed with the suppliers.

All performance information was determined in a controlled environment. Actual results may vary. Performance information is provided "AS IS" and no warranties or guarantees are expressed or implied by IBM.

The IBM home page on the Internet can be found at <http://www.ibm.com>.

The IBM BladeCenter JS20 home page on the Internet can be found at <http://www.ibm.com/servers/eserver/bladecenter/index.html>.