



Installing SUSE LINUX Enterprise Server 9 SP1 on an IBM @server BladeCenter JS20

September 2005

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

Contributors:
Noel Santiago-Ramos
IBM Corporation
noelsr@us.ibm.com

Aaron T. Bolding
IBM Corporation
atboldin@us.ibm.com



I. Introduction.....	3
Prerequisites.....	3
IP address scheme	4
Remote access to the JS20.....	4
II. Manual installation from CD.....	6
Setup to install from CD	7
Start the installation.....	7
Install using the YaST screens	8
III. Manual installation over network	9
Setup of the install server.....	9
Reusing a SLES9 installation tree	15
Setup to install over network	16
Start the installation.....	16
Define installation settings using linuxrc	17
Install using the YaST screens	18
IV. Auto-installation over network via OpenFirmware	19
Setup of the install server.....	19
Install using OpenFirmware.....	21
V. Auto-installation over network, patch install kernel	22
Setup of the install server.....	23
Setup to install over network	23
Install using patched kernel.....	24
VI. Auto-installation over network using CSM.....	25
CSM management server (MS) setup.....	25
Define the Nodes to the MS	29
CSM SLES9 SP1 installation setup	33
Install the OS to the nodes	35
Monitor and verify the installation	35
Monitoring with CSM - optional.....	37
VII. YaST Online Update.....	38
YOU Server setup.....	39
YOU Client setup	39
VIII. Alternatives	39
VNC or SSH.....	39
Power control of the JS20	39
IX. Troubleshooting	40
Serial over LAN problems	40
DHCP/BOOTP problems.....	40
TFTP problems	40
CSM	41
X. References.....	42



I. Introduction

This whitepaper explains how to install SUSE LINUX® Enterprise Server 9 Service Pack 1 (SP1) onto an IBM @server® BladeCenter™ JS20 blade using a variety of install procedures. The install procedures, their incremental differences and typical use case are as follows:

- Manual installation from CD
 - no install server, step through the YaST screens
 - use when there are few systems, infrequent installs
- Manual installation over network
 - create install server, step through the YaST screens
 - use when installing a few more systems with varying configurations
- Auto-installation over network via OpenFirmware
 - create install server and AutoYaST control file
 - use when:
 - frequently installing many systems with a similar configuration or categories of configurations
 - controlling the kernel options from each install target
- Auto-installation over network, patching install kernel
 - create install server, AutoYaST control file and patch kernel options on install server
 - use when:
 - frequently installing many systems with a similar configuration or categories of configurations
 - controlling the kernel options from the install server by patching the install kernel
- Auto-installation over network using Cluster Systems Management (CSM)
 - CSM creates the install server, AutoYaST control file and handles the kernel options
 - use when:
 - full function, yet easy-to-use distributed system management solution is necessary for 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 AutoYaST configuration 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

Serial over LAN Setup

This document assumes the Serial over LAN feature is already setup on the BladeCenter. 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 function of the blade server. To find out more about Serial over LAN setup, refer to the *Serial over LAN Setup Guide – IBM eServer™ BladeCenter and BladeCenter T* document.



One ESM or two ESMs

The systems used to generate the following install scenarios have two Ethernet Switch Modules (ESM). Therefore I performed the installs over ESM #2 (Bay 2, lower bay) to minimize the disruption of traffic on ESM #1 (Bay 1, upper bay), which is carrying Serial over LAN traffic. If only one ESM is in place, the installation and Serial over LAN traffic have to coexist on the same external ESM interface. This may cause Serial over LAN session drops. But the session can always be restarted.

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, hostname = ms:

- IP address is 192.168.70.50/255.255.255.0

BladeCenter components:

- management module, hostname = mm:
 - external is static IP at 192.168.70.125/255.255.255.0
 - to blades is static IP at 192.168.70.126/255.255.255.0
 - to IO modules is static IP at 192.168.70.127/255.255.255.0
- ethernet switch module 1, hostname = esm1:
 - static IP at 192.168.70.130/255.255.255.0
- ethernet switch module 2, hostname = esm2:
 - static IP at 192.168.70.131/255.255.255.0

Blade1, hostname = blade1:

- eth0 is disabled, to accommodate only SOL usage
- eth1 is dynamic IP at 192.168.70.100/255.255.255.0
(IP assigned to MAC address in dhcpd.conf)

Blade2, hostname = blade2:

- eth0 is disabled, to accommodate only SOL usage
- eth1 is dynamic IP at 192.168.70.101/255.255.255.0
(IP assigned to MAC address in dhcpd.conf)

Gateway:

- IP address is 192.168.70.1

Remote access to the JS20

This section describes mechanisms for accessing and controlling the JS20.

1. serial console access - Serial over LAN (SOL) session
2. GUI-based installation
 - a. Virtual Network Computing (VNC) session
 - b. Secure Shell (SSH) session
3. hardware control - management module's (MM) web interface



Using a Serial over LAN session

The Serial over LAN interface enables a JS20 to be directly managed from a remote location by providing serial console access. 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/PASSWORD (note: all upper case, with a zero). This takes you to the MM command line interface (CLI). To bring up a console for a blade, enter this command at the CLI prompt:

```
system> console -T 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 (eg. blade in slot number 2), enter:

```
system> env -T system:blade[2]  
OK  
system:blade[2]>
```

Then one just needs to enter:

```
system:blade[2]> console
```

to start an SOL session on blade x.

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

Using a Virtual Network Computing session

If requested, SLES installation will start a VNC server process during installation, on the install target. A VNC client (eg. a ThinkPad, etc...) can connect to the VNC server process in order to perform a GUI-based installation.

Installing the VNC client

To install a VNC client on a Windows® system, copy the tightvnc*setup.exe file (eg. tightvnc-1.2.9-setup.exe) from the /dosutils/tightvnc directory of the SLES9 GA CD #1. Run tightvnc*setup.exe to install the TightVNC client package.

To install a VNC client on a Linux system, download the vnc RPM package. Install the package using the YaST “Install/Remove Software” option or use the rpm command.



Launching the VNC server on the install target during installation

There are numerous ways to start the VNC server on the install target. The installation procedures in this document will utilize the VNC session whenever a GUI-based installation is desirable. To be specific, to setup a VNC server to run on the target install system, options must be passed to the kernel. These options (vnc and vncpassword) are passed to the kernel using the "install" command.

Starting the VNC client

Later, during the installation process, YaST will prompt you start the VNC client in order to connect to the VNC server. YaST will provide the server IP and display information at the time, in a message like:

"You can connect to 192.168.70.100, display :1 now with vncviewer".

At this time, start the VNC client and indicate you want to connect to VNC server "192.168.70.100:1". To start the VNC client on Windows, click on the "TightVNC Viewer" icon (vncviewer.exe). To start the VNC client on Linux, run the program vncviewer. Enter the VNC password when prompted. Next, you should see the YaST screen prompting you to proceed with the installation.

Using an SSH session

SSH can be used wherever VNC is used, however in this document we've used VNC throughout the examples. For more information about using SSH to drive the installs, see http://portal.suse.com/sdb/en/2002/12/remote_install.html

Using the Management Module's Web interface

You'll often need to use the BladeCenter Management Module (MM) 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 (eg. <http://192.168.70.125> or <http://mm.bclab.ibm.com>). 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.

II. Manual installation from CD

This section describes the process of installation of SLES9 SP1 from 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
- Start the installation
- Install using the YaST screens



Setup to install from CD

Start the MM Web interface, as described in *Section 1, 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 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. You can verify that the Media Tray (MT) has been assigned to the correct blade, under "Monitors", select "System Status", then "Blade Servers". Ensure the MT column contains an X on the appropriate Blade Server. Insert the SLES9 SP1 CD1 into the media tray's CD-ROM drive.

Start the installation

In this example, we'll observe the serial console output via an SOL session, and then use the VNC client GUI capability to more easily maneuver through the YaST screens. Then we'll return to the SOL session after YaST has completed. Start an SOL session, as described in *Section 1, Using a Serial over LAN session*, in order to watch the serial console output.

To kick off the installation, restart the server, 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".

On the SOL console, you should see the progress codes scrolling past while the JS20 initializes.

When yaboot presents the "boot:" prompt, we need to supply some kernel options that indicate to start the VNC server. These parameters will be passed to the kernel. At the "boot:" prompt, enter "install vnc=1 vncpassword=PASSWORD" (just a sample password).

In a few minutes, you'll be prompted to "Make sure that CD number 1 is in your drive". Insert the SLES9 GA CD1 in the drive and specify OK at the prompt.

Next, select the language you want to use. Then, Linuxrc will start and present its main menu. Choose "Start Installation or System", then "Start Installation or Update". When prompted to "Choose the source medium", select the CD-ROM. Next, some networking selections will be presented in order to connect the VNC client and server. When asked to "Choose the network device", select eth1. eth0 carries SOL traffic and may not be connected to the switch. When asked "Automatic Configuration via DHCP", press 1 (yes) if a DHCP server is present and ready to use. Otherwise, press 0 (no) and enter the IP information.



Install using the YaST screens

In a few minutes, the message “starting VNC server...” will be displayed. There will also be connect instructions, for example “You can connect to 192.168.70.100, display :1 now with vncviewer”. Start a VNC session, as described in *Section 1, Using a Virtual Network Computing session*, in order to start the GUI-based installation.

Next, you should see the YaST screen prompting you to select a language. Select the appropriate language. At this point, YaST will begin analyzing the system. Upon completion, if the install target already has an OS installed, you’ll get prompted for a “New installation”, “Update an existing system”, “Repair an installed system”, “Boot installed system” or “Abort installation. Choose “New installation” to replace any existing installation.

On the next screen you will be shown the installation settings for keyboard, mouse, partitioning, software, boot, time zone and language. For example, click on “Change”, then select “Software” if you want to choose a minimum configuration or additional software packages.

Note, on partitioning, if the system has been previously installed follow this procedure to ensure a bootable partitioning scheme is created.

1. Click on “Change”, and then select “Partitioning”.
2. Click on “Create custom partition setup”, then “Next”
3. Choose “/dev/hda” from the list of hard disks, then “Next”
4. Choose “Use Entire Hard Disk”, then “Next”
5. returns to Installation Settings window

When finished with all the Installation Settings, click on Accept.

At this point you’ll be given one last warning before proceeding with the installation. To proceed, select “Yes, install”. When prompted, insert the correct SLES 9 GA or SP CD.

Now you’ll see a panel containing 3 sub-panels: Current Package, Installation and Installation Log. The Installation sub-panel shows the progress of the overall installation process in the upper right hand corner.

After about 18-20 minutes, after all the packages are installed, basic installation is finished. YaST will update the configuration, copy files to the installed system, install the boot manager and prepare the system for the initial boot.

During the initial boot, enter “linux” at the boot prompt, or let it timeout. Prior to presenting the login prompt, the message “starting VNC server...” will be displayed. This is done in order to finish the installation. Start a VNC session, as described in *Section 1, Using a Virtual Network Computing session*, in order to use an X-server screen to finish the installation.

On the X-server screen, you should be asked to confirm hardware detection, press Continue. Next, you should see the YaST screen prompting you to specify the password for the “root” user. After that you can also provide further configuration for the: Network, Online Update, Services and Users. After configuring the Network, you can “Test the Internet Connection”. This step allows you to check for the latest updates to SLES9 SP1. Upon completion of the configuration,



YaST will save the configuration. Next, the Release Notes are presented for your viewing. Click on Next, when done reviewing the Release Notes.

Finally, YaST opens a dialog to configure the printers, graphics card and other devices. Click on a component to start its configuration. After the configuration data has been saved, click on "Finish". The system should continue booting and present a login prompt back on the SOL session. Log in to the system.

SLES9 SP1 will automatically insert the hard drive as the first device in the boot sequence. Finally, we want to reboot the installed server from the hard drive where we just performed the install. On the SOL session, enter the command:

```
# shutdown -r now
```

The server should reboot successfully from the hard drive, with network connectivity. The installation is 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 YaST screens for each install. This install procedure requires:

- Setup of the install server
 1. DHCP/BOOTP configuration
 2. TFTP configuration
 3. NFS configuration
 4. Building the installation tree
- Setup to install over network
- Start the installation
- Define installation settings
- Install using the YaST screens

Setup of the install server

In general, any system or combination of systems that provide DHCP/BOOTP, TFTP, NFS and the install tree satisfy the requirements for a network-based install server.

DHCP/BOOTP Configuration

First, the install server needs to be setup as a DHCP/BOOTP server. The DHCP server package that comes with SLES9 needs to have its system configuration file modified to work with JS20's. The original /etc/sysconfig/dhcp should contain a line like this:

```
DHCPD_BINARY=""
```

To install JS20's, this line needs to be changed to:



```
DHCPD_BINARY="/usr/sbin/dhcpd.lpf"
```

Next, create a /etc/dhcpd.conf file, here's an example.

```
# dhcpd.conf
#
# Sample configuration file for ISC dhcpd
#

# 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 DHCP client 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. By directing DHCP/BOOTP through the JS20's eth1 interface, SOL via the MM will function uninterrupted on eth0.

Some issues to note about dhcpd.conf:

- Most tftp daemons change the root directory (using the chroot command) to something different than /, usually /tftpboot. This is dictated by 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).

Note, you can use the command `rcdhcpd` to control the execution of the DHCP server daemon (`dhcpd`).

```
# rcdhcpd
```

```
Usage: /usr/sbin/rcdhcpd {start|stop|status|try-restart|restart|force-reload|reload|probe|check-syntax} [-v]
```

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

```
# rcdhcpd restart
Shutting down DHCP server           done
Starting DHCP server [chroot]       done
```

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

```
#
# hosts          This file describes a number of hostname-to-address
#                mappings for the TCP/IP subsystem.  It is mostly
#                used at boot time, when no name servers are running.
#                On small systems, this file can be used instead of a
#                "named" name server.
# Syntax:
#
# IP-Address    Full-Qualified-Hostname  Short-Hostname
#
127.0.0.1      localhost

# special IPv6 addresses
::1           localhost ipv6-localhost ipv6-loopback

fe00::0       ipv6-localnet

ff00::0       ipv6-mcastprefix
ff02::1       ipv6-allnodes
ff02::2       ipv6-allrouters
ff02::3       ipv6-allhosts

192.168.70.50  ms.bclab.ibm.com    ms
192.168.70.125 mm.bclab.ibm.com    mm
192.168.70.130 esm1.bclab.ibm.com  esm1
```



```
192.168.70.131 esm2.bclab.ibm.com esm2

192.168.70.100 blade1.bclab.ibm.com blade1
192.168.70.101 blade2.bclab.ibm.com blade2

# end of hosts
```

TFTP Configuration

Next we need to setup the install server as a TFTP server. To do this, start YaST by running “yast” from the command line. Then in “Network Services” select “TFTP Server”. At this point, if TFTP is not installed, you will be prompted to install the xinetd and tftp packages. Click on Continue to install those packages. Then, on the TFTP configuration screen, select Enable and specify a Boot Image Directory of /tftpboot. Then click Finish and reply Yes if prompted to create the Boot Image Directory.

Note, if “TFTP Server” is not an available option in “Network Services”, click on “Network Services (inetd)”. In the “Currently Available Services” window, scroll down to the “tftp” service entry. Click on Edit, and ensure the TFTP “Service is active” and the server arguments indicate “-s /tftpboot”.

Then we need to copy the */install* file from the SLES9 SP1 CD #1 to the /tftpboot directory.

Note, since the tftpd process usually does NOT run as the root user, we need to ensure the files in /tftpboot have the correct owner and group IDs assigned. Verify by starting YaST again, run “yast” from the command line. In “Network Services” select “Network Services (inetd)”. In the “Currently Available Services” window, scroll down to the “tftp” service entry. Click on Edit, check the User and Group value. In our case, a User value and Group value of “tftpd” is specified. The user and group IDs of the files in /tftpboot must match these User and Group values (ie. tftpd). To change the file to match the User and Group values, run this:

```
# chown tftpd:tftpd /tftpboot/install
```

Alternatively, one could alter the User value and Group value in the TFTP Service entry to match the files in /tftpboot.

NFS Configuration

Next, we need to setup the install server as an NFS server. To do this, start YaST by running “yast” from the command line. Then in “Network Services” select “NFS server”. Select “Start NFS server” and select Next. Under “Directories to export to others”, click on “Add directory”, specify “/install-tree” and click OK. Let YaST create the directory, if necessary. When prompted for the Hosts wildcard, specify “*” and for Options, specify “ro,root_squash,sync”. Note, those are the defaults presented by YaST for Hosts wildcard and Options. Click on OK for the Hosts wildcard and Options. Then click on Finish to complete the NFS configuration.

Note, you can also use the command rcnfs to control the execution of the NFS daemon (nfsd).

```
# rcnfs
Usage: /usr/sbin/rcnfs {start|stop|status|reload|force-reload|restart|try-restart}
```



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. In a flat installation tree, all the CDs are copied to one directory and the package database is rebuilt. For a multiple source installation tree, the CDs are copied into separate directories and control files are created. We'll create a multiple source installation tree.

Note, if you already have an installation tree built for SLES9, or if you intend to serve out installs for SLES9 and SLES9 SP1 separately, read this section and the next section before proceeding. You'll be able to save some space on your install server, should that be an issue.

The installation tree will be built in the `/install-tree/sles91` directory. Note, the `/install-tree/sles91` directory name is used to accommodate a hierarchy for serving other distributions (ie. SLES8 in `/install-tree/sles8`, SLES9 in `/install-tree/sles9`). If you're only installing SLES9 SP1, 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 SLES9 GA CDs to disk. Be sure to specify the "periods" in the commands that follow. Also note, CD #2 is copied to `/install-tree/sles91/core/CD1`, and the remaining 4 CDs follow the same pattern. This in effect rennumbers the last 5 CDs.

```
# mkdir -p /install-tree/sles91/sles/CD1 ; cd /install-tree/sles91/sles/CD1
# umount /media/cdrom
```

insert SLES9 GA CD #1

```
# mount /dev/cdrom /media/cdrom ; cp -a /media/cdrom/* . ; umount /media/cdrom
# mkdir -p /install-tree/sles91/core/CD1 ; cd /install-tree/sles91/core/CD1
```

insert SLES9 GA CD #2

```
# mount /dev/cdrom /media/cdrom ; cp -a /media/cdrom/* . ; umount /media/cdrom
# mkdir -p /install-tree/sles91/core/CD2 ; cd /install-tree/sles91/core/CD2
```

insert SLES9 GA CD #3

```
# mount /dev/cdrom /media/cdrom ; cp -a /media/cdrom/* . ; umount /media/cdrom
# mkdir -p /install-tree/sles91/core/CD3 ; cd /install-tree/sles91/core/CD3
```

insert SLES9 GA CD #4

```
# mount /dev/cdrom /media/cdrom ; cp -a /media/cdrom/* . ; umount /media/cdrom
# mkdir -p /install-tree/sles91/core/CD4 ; cd /install-tree/sles91/core/CD4
```

insert SLES9 GA CD #5



```
# mount /dev/cdrom /media/cdrom ; cp -a /media/cdrom/* . ; umount /media/cdrom  
# mkdir -p /install-tree/sles91/core/CD5 ; cd /install-tree/sles91/core/CD5
```

```
insert SLES9 GA CD #6
```

```
# mount /dev/cdrom /media/cdrom ; cp -a /media/cdrom/* . ; umount /media/cdrom
```

Next, copy the content of the SLES9 SP1 CDs to disk. Again, be sure to specify the “periods” in the commands that follow.

```
# mkdir -p /install-tree/sles91/sp1/CD1 ; cd /install-tree/sles91/sp1/CD1  
# umount /media/cdrom
```

```
insert SLES9 SP1 CD #1
```

```
# mount /dev/cdrom /media/cdrom ; cp -a /media/cdrom/* . ; umount /media/cdrom  
# mkdir -p /install-tree/sles91/sp1/CD2 ; cd /install-tree/sles91/sp1/CD2
```

```
insert SLES9 SP1 CD #2
```

```
# mount /dev/cdrom /media/cdrom ; cp -a /media/cdrom/* . ; umount /media/cdrom  
# mkdir -p /install-tree/sles91/sp1/CD3 ; cd /install-tree/sles91/sp1/CD3
```

```
insert SLES9 SP1 CD #3
```

```
# mount /dev/cdrom /media/cdrom ; cp -a /media/cdrom/* . ; umount /media/cdrom
```

Link in some files from the CD images. Again, be sure to specify the ending “periods” in these commands.

```
# cd /install-tree/sles91  
# cp -l -a sles/CD1/boot .  
# cp -l -a sles/CD1/content .  
# cp -l -a sles/CD1/control.xml .  
# cp -l -a sles/CD1/media.1/ .  
# cp -l -a sp1/CD1/driverupdate .  
# cp -l -a sp1/CD1/linux/ .  
# mkdir yast  
# mkdir yastcfg
```

Note, the yastcfg directory will be used in the auto-installation processes. Auto-installation over the network shares most of the same install server setup as manual installation over the network.

Create the order and instorder files. The order file is a tab-delimited file that pairs a directory (ie. the location of the description) with an installation source (ie. the location of the product, defaults to /).

```
# printf "sp1/CD1\tsp1/CD1\n" > yast/order  
# printf "sles/CD1\tsles/CD1\n" >> yast/order  
# printf "core/CD1\tcore/CD1\n">> yast/order
```



The instorder file indicates the order in which to install the installation sources.

```
# printf "sp1/CD1\tsp1/CD1\n" > yast/instorder
# printf "sles/CD1\tsles/CD1\n" >> yast/instorder
# printf "core/CD1\tcore/CD1\n">> yast/instorder
```

Verify that order and instorder look like this:

```
# more yast/order
sp1/CD1 sp1/CD1/
sles/CD1      sles/CD1/
core/CD1      core/CD1/

# more yast/instorder
sp1/CD1 sp1/CD1/
sles/CD1      sles/CD1/
core/CD1      core/CD1/
```

That completes the building of the installation tree. The directory layout of the installation tree should look like this:

```
/install-tree/
  sles91/
    sles/
      CD1/
    core/
      CD1/, CD2/, CD3/, CD4/, CD5/
    sp1/
      CD1/
    boot/
    content
    control.xml
    driverupdate
    linux
    media.1/
    yast/
      order
      instorder
    yastcfg/
```

Reusing a SLES9 installation tree

If you already have an installation tree for SLES9 that looks like this:

```
/install-tree/
  sles9/
    sles/
      CD1/
    core/
      CD1/, CD2/, CD3/, CD4/, CD5/
```



```
boot/  
content  
control.xml  
media.1/  
yast/  
  order  
  instorder  
yastcfg/
```

You can create an installation tree, like the one below, for SLES9 SP1 which contains only the SP1 CDs and reuse the SLES9 installation tree.

```
/install-tree/  
  sles91/  
    sles, link to /install-tree/sles9/sles  
    core, link to /install-tree/sles9/core  
    sp1/  
      CD1/  
        boot/  
        content  
        control.xml  
        driverupdate  
        linux  
        media.1/  
        yast/  
          order  
          instorder  
        yastcfg/
```

This allows you to continue to separately serve SLES9 and SLES9 SP1 installs, since there are control files in place for both installation trees. It also saves disk space on the install server.

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

In this example, we'll observe the serial console output via an SOL session, and then use the VNC client GUI capability to more easily maneuver through the YaST screens. Start an SOL session, as described in *Section I, Using a Serial over LAN session*, in order to watch the serial console output.



To kick off the installation, restart the server, 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”.

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 command line 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 60-90 seconds until the SOL session can be restarted.

Define installation settings using linuxrc

linuxrc is a manual setup program that runs the start-up stage of the kernel prior to the actual boot process. It can be used to define installation settings. We'll also use it to setup VNC. Upon completion, linuxrc passes control to YaST which performs the installation. At the linuxrc prompt, “>”, enter these options, followed by the Enter key:

- select your language, press 4 for English
- press 1 for “Settings”
- press 4 for “Debug (expert)”
- press 8 for “VNC Enable or Disable”
- press 1 for “Yes”
- continue pressing Enter until back at the main menu

That concludes the setup for VNC. Since we're booting from the network we need to load the kernel module for the network/Ethernet adapter.

- press 3 for “Kernel Modules (Hardware Drivers)”
- press 1 for “Load ppc_pseries64 Modules”
- press 5 for “bcm5700 : Broadcom BCM5700”
- press Enter, when prompted for parameters for bcm5700
- at this time it will load the bcm5700 kernel module
- when finished, press Enter until back at the main menu

That concludes the loading of the Ethernet kernel module. Next we'll start the installation, provide the VNC password and some networking information.

- press 4 for “Start Installation of System”
- press 1 for “Start Installation or Update”
- press 2 for “Network”
- press 3 for “NFS”
- provide the VNC password, for example PASSWORD
- press 2 for “eth1”
- press 1 for “Yes” (indicating Automatic configuration via DHCP)
- supply the IP address of the NFS server (192.168.70.50), when requested
- supply the directory path (/install-tree/sles91), when requested

That completes linuxrc, the kernel will continue to boot up and pass control to YaST for the installation.



Install using the YaST screens

In a few seconds, the message “starting VNC server...” will be displayed. There will also be connect instructions, for example “You can connect to 192.168.70.100, display :1 now with vncviewer”. Start a VNC session, as described in *Section 1, Using a Virtual Network Computing session*, in order to start the GUI-based installation.

Next, you should see the YaST screen prompting you to accept the license. Then you’ll need to select a language. Select the appropriate language. If the install target already has an OS installed, you’ll get prompted for a “New installation”, “Update an existing system”, “Repair an installed system”, “Boot installed system” or “Abort installation. Choose “New installation” to replace any existing installation.

On the next screen you will be shown the installation settings for keyboard, mouse, partitioning, software, boot, time zone and language. For example, click on “Change”, then select “Software” if you want to choose a minimum configuration or additional software packages.

Note, on partitioning, if the system has been previously installed follow this procedure to ensure a bootable partitioning scheme is created.

1. Click on “Change”, then select “Partitioning”.
2. Click on “Create custom partition setup”, then “Next”
3. Choose “/dev/hda” from the list of hard disks, then “Next”
4. Choose “Use Entire Hard Disk”, then “Next”
5. returns to Installation Settings window

When finished with all the Installation Settings, click on Accept.

At this point you’ll be given one last warning before proceeding with the installation. To proceed, select “Yes, install”.

Now you’ll see a panel containing 3 sub-panels: Current Package, Installation and Installation Log. The Installation sub-panel shows the progress of the overall installation process in the upper right hand corner.

After about 12-15 minutes, after all the packages are installed, YaST boots into the newly installed system.

Prior to presenting the login prompt, the message “starting VNC server...” will be displayed. This is done in order to finish the installation. Start a VNC session, as described in *Section 1, Using a Virtual Network Computing session*, in order to use an X-server screen to finish the installation.

On the X-server screen, you should be asked to confirm hardware detection, press Continue. Next, you should see the YaST screen prompting you to specify the password for the “root” user. After that you can also provide further configuration for the: Network, Online Update, Services and Users. After configuring the Network, you can “Test the Internet Connection”. This step allows you to check for the latest updates to SLES9 SP1. Upon completion of the configuration, YaST will save the configuration. Next, the Release Notes are presented for your viewing. Click on Next, when done reviewing the Release Notes.



Finally, YaST opens a dialog to configure the printers, graphics card and other devices. Click on a component to start its configuration. After the configuration data has been saved, click on "Finish". The system should continue booting and present a login prompt back on the SOL session. Log in to the system.

SLES9 SP1 will automatically insert the hard drive as the first device in the boot sequence. Finally, we want to reboot the installed server from the hard drive where we just performed the install. On the SOL session, enter the command:

```
# shutdown -r now
```

The server should reboot successfully from the hard drive, with network connectivity. The installation is 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 (ie. utilizing the AutoYaST control 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/BOOTP configuration
 2. TFTP configuration
 3. NFS configuration
 4. Building the installation tree
 5. Create the AutoYaST control file
- Install using OpenFirmware

Setup of the install server

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

Create the AutoYaST control file

Next, we need to create an AutoYaST control file which will contain the predefined "answers" to the configuration options and be used in lieu of manually stepping through the YaST screens. These are suggested alternatives for generating the control file:

1. Create an AutoYaST control file from scratch
2. Use an existing system as a reference to create the control file
3. Modify a SLES8 or SLES9 AutoYaST control file
4. A sample JS20 AutoYaST control file is available for download from the same location that you downloaded this document from.



To create an AutoYaST control file from scratch, start YaST by running “yast” from the command line. Then in “Misc” select “Autoinstallation” (or run “yast autoyast” from the command line). From this point, a number of options are available via the toolbar, options such as:

- creating or modifying a configuration – by selecting modules from the left-hand side hierarchy and clicking on “configure”.
- cloning the install server configuration – by selecting the Tools pulldown and then Create Reference Profile.
- importing foreign configurations (eg. Alice or Kickstart configurations) – by selecting the File pulldown and then Import.

For further details on creating/modifying/customizing a control file or creating a rules file with classes, refer the documentation provided by SUSE LINUX at:
<http://www.SUSE.com/~nashif/autoinstall/9.0/html/index.html>

It's simplest to create a control file by using another system as a reference. But its configuration (ie. hard drive, partitioning, network interfaces, etc...) may not match the desired configuration of the JS20. Using an installed JS20 blade as a reference will provide the most accurate control file. Either way, it may be necessary to modify the control file (see <http://www.SUSE.com/~nashif/autoinstall/9.0/html/index.html> for details). To use an existing SLES9 installation as a reference, click on the Tools pulldown, then “Create Reference Profile”. Select among the Basic Resources and Additional Resources you'd like included in the profile. When done, click on “Create”. To save, select the File pulldown and Save or Save As. Transfer this file to the /install-tree/sles9/yastcfg directory on the install server. Note, we made the yastcfg directory earlier when building the install-tree. When done saving, click on Finish and Close yast.

Note, when generating an AutoYaST control file from a SLES8 system, to be used to install SLES9 servers, the <bootloader> section must be removed from the SLES8-based control file.

A sample JS20 AutoYaST control file is available for download from the same location that you downloaded this document from. Save this file to the /install-tree/sles91/yastcfg directory, use filename of “default”, on the install server.

If the configuration of the systems to be installed is different, the coupling of control files to install targets is accomplished by using the hexadecimal value of the node's IP address. For example, blade1's IP is 192.168.70.100, its unique AutoYaST control file would have a file name of “C0A84664”. The unique AutoYaST control file for blade2 (192.168.70.101) would have a file name of “C0A84665”. In order to avoid replicating the file for each group of similarly configured servers, a common set of hexadecimal values can be specified as a filename. For example, a control file named “C0A846” would be used by all the install targets having an IP address of 192.168.70/24. In general, a control file that matches an install target will be scanned for by truncating one hexadecimal digit at a time. Therefore, the search order for blade1's AutoYaST control file will be:

C0A84664
C0A8466
C0A846
C0A84
C0A8
C0A
C0
C



```
000D601E0E8D
default
```

The next to last entry in the search order is the install target's MAC address.

Note, in the syslinux RPM package there exists a utility called gethostip, which will generate the hexadecimal value of an IP address.

Install using OpenFirmware

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

In this example, we'll observe the serial console output via an SOL session. For the auto-installation procedures we won't start a VNC session, since there's nothing interactive occurring. Start an SOL session, as described in *Section 1, Using a Serial over LAN session*, in order to watch the serial console output.

To kick off the installation, restart the server, 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".

On the SOL console, you should see the progress codes scrolling past while the JS20 initializes. When installing via OpenFirmware, the SOL session should not be terminated by the install process, as it was when installing over the network.

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 "0 >" prompt. At the OpenFirmware prompt, run the "ls" command, the output should contain a /pci section that looks like:

```
0 > ls
.....
000000d34278: /pci@800000f8000000
000000d37d28: /pci@0
000000d4b408: /ethernet@1
000000d5a788: /ethernet@1,1
.....
ok
```

We need to issue the following "boot" command:

```
0 > boot /pci@800000f8000000/pci@0/ethernet@1,1:bootp,192.168.70.50,,192.168.70.100,0.0.0.0
autoyast=nfs://192.168.70.50/install-tree/sles91/yastcfg/ install=nfs://192.168.70.50/install-tree/sles91
netdevice=eth1
```

Where each of the fields is defined as:

- full path name of eth1 device = /pci@800000f8000000/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 = 0.0.0.0 (do not need to send bootp request off subnet)
- arguments passed to the install kernel are:
 - location of AutoYaST control file = “autoyast=nfs://192.168.70.50/install-tree/sles91/yastcfg”
 - location of install tree = “install=nfs://192.168.70.50/install-tree/sles91”
 - directive for JS20 to install over eth1 = netdevice=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>.

When the installation starts, there will be a screen shown to accept the license. If it isn't accepted or declined manually within 10 seconds or so, it is automatically accepted and the installation proceeds. The remainder of the installation process will be visible via the SOL session. Once the installation is complete, a “login:” prompt will be displayed on the SOL session. Unless you performed customizations of the AutoYaST control file, the network interfaces will not be defined. To define the network interfaces, login to the server as root. The password will be the password specified in the AutoYaST control file. If you cloned a system to create the control file, it will be the password of the cloned system.

To define the network interfaces, start YaST by running “yast” from the command line. Then in “Network Devices” select “Network Card”. Then select “Configure” while highlighting the Network Card with the highest MAC address, that will be the eth1 device. Then configure the IP settings: DHCP/static, Hostname and name server, Routing and any DHCP client options. Now press “Finish” and “Quit” out of YaST.

SLES9 SP1 will automatically insert the hard drive as the first device in the boot sequence. Finally, we want to reboot the installed server from the hard drive where we just performed the install. On the SOL session, enter the command:

```
# shutdown -r now
```

The server should reboot successfully from the hard drive, with network connectivity. The installation is complete.

V. Auto-installation over network, patch install kernel

This install procedure is best used when frequently installing many systems with a similar configuration or categories of configurations. In this procedure, we'll patch the install kernel on the install server to contain kernel options. This install procedure requires:

- Setup of the install server
 1. DHCP/BOOTP configuration
 2. TFTP configuration



3. NFS configuration
 4. Building the installation tree
 5. Create the AutoYaST control file
 6. Patch kernel options on the install server
- Setup to install over network
 - Install using patched kernel

Setup of the install server

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

Create the AutoYaST control file just as was done in *Section IV, Create the AutoYaST control file*.

Patch kernel options on the install server

The final piece of setup on the install server is to insert some kernel options into a patch area within the install kernel. These kernel options will be given to the install kernel when it starts. Previously, we copied the */install* file from the SLES9 SP1 CD #1 to the */tftpboot* directory. The kernel options to be supplied will inform each install target:

1. where its AutoYaST control file is located
2. where the installation tree is located
3. to use eth1 for the network install (since SOL is taking place over eth0)

This is done using a utility that comes with SLES9 SP1 called *mkzimage_cmdline*. This utility is located in the */ppc/netboot* directory of SLES9 SP1 CD #1. Based on the install server setup we've done, the *mkzimage_cmdline* we would invoke would be:

```
mkzimage_cmdline -s "autoyast=nfs://192.168.70.50/install-tree/sles91/yastcfg/  
install=nfs://192.168.70.50/install-tree/sles91 netdevice=eth1" -a 1 install
```

Where the general syntax for *mkzimage_cmdline* is:

```
mkzimage_cmdline filename  
work with zImage named 'filename'  
[-h] display this help  
[-v] display version  
[-a 0|1] disable/enable built-in cmdline  
overrides whatever is passed from OpenFirmware  
[-s STRING] store STRING in the boot image  
[-c] clear previous content before update  
no option will show the current settings in 'filename'
```

Again, there are many options that can be set other than *autoyast=*, *install=* and *netdevice=*. For a complete list of options, consult the SUSE LINUX documentation at <http://www.SUSE.com/~nashif/autoinstall/9.0/html/index.html>

Setup to install over network



Start the MM Web interface, as described in *Section 1, 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.

Install using patched kernel

In this example, we'll observe the serial console output via an SOL session. For the auto-installation procedures we won't start a VNC session, since there's nothing interactive occurring. Start an SOL session, as described in *Section 1, Using a Serial over LAN session*, in order to watch the serial console output.

To kick off the installation, restart the server, 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".

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 command line 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 60-90 seconds until the SOL session can be restarted.

When the installation starts, there will be a screen shown to accept the license. If it isn't manually accepted or declined within 10 seconds or so, it is automatically accepted and the installation proceeds. The remainder of the installation process will be visible via the SOL session. Once the installation is complete, a "login:" prompt will be displayed on the SOL session. Unless you performed customizations of the AutoYaST control file, the network interfaces will not be defined. To define the network interfaces, login to the server as root. The password will be the password specified in the AutoYaST control file. If you cloned a system to create the control file, it will be the password of the cloned system.

To define the network interfaces, start YaST by running "yast" from the command line. Then in "Network Devices" select "Network Card". Then select "Configure" while highlighting the Network Card with the highest MAC address, that will be the eth1 device. Then configure the IP settings: DHCP/static, Hostname and name server, Routing and any DHCP client options. Now press "Finish" and "Quit" out of YaST.

SLES9 SP1 will automatically insert the hard drive as the first device in the boot sequence. Finally, we want to reboot the installed server from the hard drive where we just performed the install. On the SOL session, enter the command:

```
# shutdown -r now
```

The server should reboot successfully from the hard drive, with network connectivity. The installation is complete.



VI. Auto-installation over network using CSM

This section describes how to populate the blade servers using Cluster Systems Management (CSM) for Linux V1.4.0.12. This install option is most useful when you want the blade servers to operate as a cluster of nodes. CSM is a deployment and management application that provides 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 SLES9 SP1 (AutoYaST) installation setup
- Install the OS to the nodes
- Monitoring with CSM - optional

CSM management server (MS) setup

Relationship between OS on MS and OS on nodes

In the following CSM install procedure, the CSM MS is assumed to be running the same Linux distribution and version that is to be installed on the nodes. This is certainly not a requirement. See the *CSM for Linux V1.4 Planning and Installation Guide*, for a description of the Linux coexistence considerations. As an example, if the Linux version differs between the MS and the nodes; special consideration needs to be given when defining the nodes, since the nodes inherit many attributes from the MS.

CSM MS installation

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 host name resolution, here's an example:

```
#
# hosts          This file describes a number of hostname-to-address
#               mappings for the TCP/IP subsystem.  It is mostly
#               used at boot time, when no name servers are running.
#               On small systems, this file can be used instead of a
```



```
#           "named" name server.
# Syntax:
#
# IP-Address  Full-Qualified-Hostname  Short-Hostname
#
127.0.0.1      localhost

# special IPv6 addresses
::1           localhost ipv6-localhost ipv6-loopback

fe00::0       ipv6-localnet

ff00::0       ipv6-mcastprefix
ff02::1       ipv6-allnodes
ff02::2       ipv6-allrouters
ff02::3       ipv6-allhosts

192.168.70.50  ms.bclab.ibm.com    ms
192.168.70.125 mm.bclab.ibm.com    mm
192.168.70.130 esm1.bclab.ibm.com  esm1
192.168.70.131 esm2.bclab.ibm.com  esm2

192.168.70.100 blade1.bclab.ibm.com blade1
192.168.70.101 blade2.bclab.ibm.com blade2

# end of hosts
```

Download the Autoupdate RPM to a temporary directory, for example: /tmp/csmreqs. The Autoupdate RPM can be downloaded from: <ftp://ftp.istrata.com/pub/autoupdate/>.

You'll also need to 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-1.4.0.12.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 SLES9 GA CDs and SLES9 SP1 CDs when prompted.

When prompted about potentially conflicting TFTP packages, select option 1 to remove the existing TFTP package 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 SLES9 GA CDs and SLES9 SP1 CDs when prompted.

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 /media/cdrom  
# csmconfig -L /media/cdrom/csmlum.full
```



Verify the full license is in place by running:

```
# csmconfig
```

The output should look like this:

```
AddUnrecognizedNodes = 0 (no)
ClusterSNum =
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 =
ClusterTM = 9078-160
DeviceStatusFrequency = 12
DeviceStatusSensitivity = 8
ExpDate = Sat Apr 23 18:59:59 2005
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)
```



Note the highlighted fields in the last two output examples, 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          (lower case L)
```

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. use the `Ishwinfo` command to gather attributes from the HW control point
2. transform the `Ishwinfo` output to a node definition file format
3. make alterations to the node definition file
4. issue the `definnode` command and provide the node definition file

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 -s -o /tmp/csm/hwinfo1
```

The resulting `hwinfo1` file should look like this:

```
# Hostname::PowerMethod::HWControlPoint::HWControlNodeID::LParID::HWType::HWModel::HWSerialNum::  
DeviceType::UUID  
blade1::blade::mm.bclab.ibm.com::blade1:::8842::41X::ZJ1W6745E17J::  
::027D B780 9B44 11E8 A0C0 000D 601E 80DC  
blade1::blade::mm.bclab.ibm.com::blade2:::8842::41X::ZJ1W6745E17G::  
::0D47 A900 73E7 11C0 87CF 000D 601E 8114
```

Note, by specifying `-s` on `Ishwinfo`, the `HWControlNodeID` field is copied to the `Hostname` field. The `HWControlNodeID` field is the same as the Blade Server's name in the MM web interface. If you want the `Hostname` and `HWControlNodeID` to be different, edit the `Hostname` fields in the `hwinfo1` file.

Now we'll use CSM's `definnode` command to create a node definition file based on the `hwinfo1` file. 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 `definnode` command. Run this `definnode` command:

```
# definnode -M hwinfo1 -C mm.bclab.ibm.com::1:2 -s ConsoleMethod=blade \  
InstallAdapterName=eth1 InstallDiskType=ide ConsoleSerialDevice= > nodedef1
```

Where, each parameter's meaning is:

- `-M hwinfo1`, identifies mapping file of IP hostname to hardware control information



- -C mm.bclab.ibm.com::1:2, provide the console server name, a starting port number and a count of following ports to use.
- -s, send output to stdout only, do not issue the definenode command
- ConsoleMethod=blade, set ConsoleMethod node attribute to "blade" for each node
- InstallAdapterName=eth1, perform each node's install over eth1
- InstallDiskType=ide, each node's install will be performed on an IDE disk
- ConsoleSerialDevice=, because the system has no serial port

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

blade1.bclab.ibm.com :

```
ConsoleMethod=blade  
ConsolePortNum=1  
ConsoleSerialDevice=''  
ConsoleServerName=mm.bclab.ibm.com  
HWControlNodeId=blade1  
HWControlPoint=mm.bclab.ibm.com  
HWModel=41X  
HWSerialNum=ZJ1W6745E17J  
HWType=8842  
InstallAdapterName=eth1  
InstallAdapterType=ent  
InstallCSMVersion=1.4.0  
InstallDiskType=ide  
InstallDistributionName=SLES  
InstallDistributionVersion=9  
InstallOSName=Linux  
InstallPkgArchitecture=ppc64  
InstallServiceLevel=SP1  
ManagementServer=ms.bclab.ibm.com  
PowerMethod=blade  
UUID=027DB7809B4411E8A0C0000D601E80DC
```

blade2.bclab.ibm.com :

```
ConsoleMethod=blade  
ConsolePortNum=2  
ConsoleSerialDevice=''  
ConsoleServerName=mm.bclab.ibm.com  
HWControlNodeId=blade2  
HWControlPoint=mm.bclab.ibm.com  
HWModel=41X  
HWSerialNum=ZJ1W6745E17G  
HWType=8842  
InstallAdapterName=eth1  
InstallAdapterType=ent  
InstallCSMVersion=1.4.0  
InstallDiskType=ide  
InstallDistributionName=SLES  
InstallDistributionVersion=9  
InstallOSName=Linux  
InstallPkgArchitecture=ppc64  
InstallServiceLevel=SP1
```



```
ManagementServer=ms.bclab.ibm.com  
PowerMethod=blade  
UUID=0D47A90073E711C087CF000D601E8114
```

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

```
# definode -f nodedef1
```

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

```
# lsnode -l (lower case L)
```

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

```
Hostname = blade1.bclab.ibm.com  
AllowManageRequest = 0 (no)  
CSMVersion =  
ChangedAttributes = {ControlFlag}  
ConfigChanged = 0 (no)  
ConsoleMethod = blade  
ConsolePortNum = 1  
ConsoleSerialDevice =  
ConsoleSerialSpeed = 9600  
ConsoleServerName = mm.bclab.ibm.com  
ConsoleServerNumber =  
FWSvcProc =  
FWSysBIOS =  
HWControlNodeId = blade1  
HWControlPoint = mm.bclab.ibm.com  
HWModel = 41X  
HWSerialNum = RAZ0053  
HWType = 8842  
InstallAdapterDuplex =  
InstallAdapterGateway =  
InstallAdapterMacaddr =  
InstallAdapterName = eth1  
InstallAdapterNetmask =  
InstallAdapterSpeed =  
InstallAdapterType = ent  
InstallCSMVersion = 1.4.0  
InstallDisk =  
InstallDiskType = ide  
InstallDistributionName = SLES  
InstallDistributionVersion = 9  
InstallKernelVersion =  
InstallMethod =  
InstallOSName = Linux  
InstallPkgArchitecture = ppc64  
InstallServer =  
InstallServiceLevel = SP1  
InstallStatus = PreManaged  
InstallTemplate =
```



LParID =
LastCFMUpdateTime =
ManagementServer = ms.bclab.ibm.com
Mode = PreManaged
Name = blade1.bclab.ibm.com
NodeNameList = {ms}
PowerMethod = blade
PowerStatus = 1 (on)
Status = 1 (alive)
UUID = 027DB7809B4411E8A0C0000D601E80DC
UpdatenodeFailed = 0 (false)
UserComment =

Hostname = blade2.bclab.ibm.com
AllowManageRequest = 0 (no)
CSMVersion =
ChangedAttributes = {ControlFlag}
ConfigChanged = 0 (no)
ConsoleMethod = blade
ConsolePortNum = 2
ConsoleSerialDevice =
ConsoleSerialSpeed = 9600
ConsoleServerName = mm.bclab.ibm.com
ConsoleServerNumber =
FWSvcProc =
FWSysBIOS =
HWControlNodeId = blade2
HWControlPoint = mm.bclab.ibm.com
HWModel = 41X
HWSerialNum = RAZ0035
HWType = 8842
InstallAdapterDuplex =
InstallAdapterGateway =
InstallAdapterMacaddr =
InstallAdapterName = eth1
InstallAdapterNetmask =
InstallAdapterSpeed =
InstallAdapterType = ent
InstallCSMVersion = 1.4.0
InstallDisk =
InstallDiskType = ide
InstallDistributionName = SLES
InstallDistributionVersion = 9
InstallKernelVersion =
InstallMethod =
InstallOSName = Linux
InstallPkgArchitecture = ppc64
InstallServer =
InstallServiceLevel = SP1
InstallStatus = PreManaged
InstallTemplate =
LParID =



```
LastCFMUpdateTime =  
ManagementServer = ms.bclab.ibm.com  
Mode = PreManaged  
Name = blade2.bclab.ibm.com  
NodeNameList = {ms}  
PowerMethod = blade  
PowerStatus = 0 (off)  
Status = 1 (alive)  
UUID = 0D47A90073E711C087CF000D601E8114  
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 SLES9 SP1 installation setup

Now we need to supply the SLES9 GA and SP1 CDs to CSM so it can build the install tree and the AutoYaST control 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 1, 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.ibm.com on  
blade2.bclab.ibm.com off          (blade2 happens to be powered off)
```

If error messages result, run the hardware control probe to diagnose the hardware control subsystem:

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

CSM maintains templates for the different automated installs it performs. If you want to alter the AutoYaST control file template for the upcoming install, the template is located in `/opt/csm/install/yastcfg.SLES9-ppc64.xml`. For further details on creating/modifying/customizing a control file or creating a rules file with classes, refer the documentation provided by SUSE LINUX at: <http://www.SUSE.com/~nashif/autoinstall/9.0/html/index.html>



Next, we need to supply the SLES9 GA CDs to CSM, run the following command:

```
# copycds InstallServiceLevel=GA
```

Note, the `InstallDistributionName=SLES` and `InstallDistributionVersion=9` node attributes are used by `copycds` by default. `InstallServiceLevel=GA` is specified in order to override the `InstallServiceLevel=SP1` currently set in the node definitions.

Insert the requested SLES9 GA CD when prompted. As CSM copies the contents of the CD, the % complete is displayed in the lower left of the window.

Next, we need to supply the SLES9 SP1 CDs to CSM, run the following command:

```
# copycds
```

Note, the default node attributes in the node definitions are `InstallDistributionName=SLES`, `InstallDistributionVersion=9` and `InstallServiceLevel=SP1`.

Insert the requested SLES9 SP1 CD when prompted. CSM maintains the content of these CDs in `/csminstall/Linux/SLES/9/ppc64`.

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

```
# monitorinstall
```

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

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

Next, we need to run:

```
# csmsetupyast -x -P
```

to create an AutoYaST control file for each of the nodes in the cluster. Where, each parameter's meaning is:

- `-x`: do not copy the CDs. No need to copy the CDs since we ran `copycds`.
- `-P`: invoke the command for all nodes where `Mode=PreManaged`.

The output of the `csmsetupyast` command can be found in `/var/log/csm/csmsetupyast.log`.

A quick word about the target of CSM commands. 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 AutoYaST configuration files in
/csminstall/csm/1.4.0/autoyast.SLES9

Install the OS to the nodes

The next step is to install SLES9 SP1 and the CSM client code on the nodes, to do this run:

```
# installnode -P
```

Monitor and verify the installation

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

```
# monitorinstall
```

typical output might be:

Node	Mode	Status
blade1.bclab.ibm.com	Installing	Rebooting and Installing Node.
blade2.bclab.ibm.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.ibm.com	Managed	Installed
blade2.bclab.ibm.com	Managed	Installed

To open up a remote console to the install targets and watch for errors or monitor the progress of the install, run:

```
# rconsole -n tempblade1, tempblade2
```

To terminate an rconsole session enter the following key sequence:

- ctrl-e, followed by
- c (a left square bracket should be displayed), followed by
- . (a period)

See the rconsole man page for additional information.

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 (lower case L)
```

The output should look similar to:

```
Name = blade1
ActiveMgtScopes = 1
Architecture = ppc64
DistributionName = SuSE
DistributionVersion = 9
KernelVersion = 2.6.5-7.139-pseries64
LoadAverage = {0.01,0.1,0.08}
NodeNameList = {blade1.bclab.ibm.com}
NumOnlineProcessors = 2
NumProcessors = 2
NumUsers = 1
OSName = Linux
PctRealMemFree = 28
PctTotalPgSpFree = 100
PctTotalPgSpUsed = 0
PctTotalTimeIdle = 100
PctTotalTimeKernel = 0
PctTotalTimeUser = 0
PctTotalTimeWait = 0
RealMemSize = 471650304
TotalPgSpFree = 230452
TotalPgSpSize = 230452
UpTime = 560
VMPgInRate = 0
VMPgOutRate = 9
VMPgSpInRate = 0
VMPgSpOutRate = 0
-----
Name = blade2
ActiveMgtScopes = 1
Architecture = ppc64
DistributionName = SuSE
DistributionVersion = 9
KernelVersion = 2.6.5-7.139-pseries64
LoadAverage = {0,0.03,0.06}
NodeNameList = {blade2.bclab.ibm.com}
NumOnlineProcessors = 2
NumProcessors = 2
NumUsers = 0
OSName = Linux
PctRealMemFree = 29
PctTotalPgSpFree = 100
PctTotalPgSpUsed = 0
PctTotalTimeIdle = 100
PctTotalTimeKernel = 0
PctTotalTimeUser = 0
PctTotalTimeWait = 0
```



```
RealMemSize = 471650304
TotalPgSpFree = 230452
TotalPgSpSize = 230452
UpTime = 851
VMPgInRate = 0
VMPgOutRate = 9
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.ibm.com: Thu Mar 3 22:44:08 CST 2005
blade2.bclab.ibm.com: Thu Mar 3 02:48:52 CST 2005
```

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 associated (mkcondresp), then activated (startcondresp). 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 associated condition and response pairs can be listed by running:

```
# lscondresp
```

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

Displaying condition with response information:

Condition	Response	Node	State
"NodeFullInstallComplete"	"RunCFMToNode"	"ms"	"Active"
"NodeManaged"	"GatherSSHHostKeys"	"ms"	"Active"
"NodeFullInstallComplete"	"GatherKRB5keytabs"	"ms"	"Active"
"pSeriesNodeFullInstallComplete"	"removeArpEntries"	"ms"	"Active"
"UpdatenodeFailedStatusChange"	"UpdatenodeFailedStatusResponse"	"ms"	"Active"
"AllServiceableSwitchEvents"	"AuditLogServiceableEvents"	"ms"	"Not active"
"NodeChanged"	"rconsoleUpdateResponse"	"ms"	"Active"
"AllServiceableHardwareEvents"	"AuditLogServiceableEvents"	"ms"	"Not active"

We'll monitor the event of a blade server changing its power state (to either on or off) and log the event. To accomplish 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 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.

VII. YaST Online Update

YaST Online Update (YOU) provides a way to download and apply patches from the SUSE LINUX Maintenance Web service. The patches can be manually or automatically downloaded and applied to keep a number of servers current. Designate a local server as a YOU server (eg. blade1.bclab.ibm.com) and synchronize it with the latest patches from SUSE LINUX. Then synchronize the remaining systems with your local YOU server.



YOU Server setup

Start YaST by running “yast” from the command line. Highlight “Software”, and then click on “Online Update”. The “Manually Select Patches” option is the default. This allows you to see the package updates available within each system component and select the packages you wish to retrieve. Click on “Next” and provide a username and password. Also, click on the checkbox to “Keep Authentication Data”, then click on “Login”.

After the desired patches have been downloaded, return to the main YaST screen, highlight “Software”, then click on “YOU Server Configuration”. Select “Start Server”, once the server has started click on “Synchronize now” to apply the patches. Note, you can also “Setup automatic synchronization” that will retrieve updates on a regularly scheduled basis.

YOU Client setup

On the YOU client side, start YaST by running “yast” from the command line. Then in “Software”, select “Online Update”. Select “New Server”, then choose a protocol such as HTTP or NFS to pull the files from the local server. Next you’ll be prompted for the Server Name (eg. blade1.bclab.ibm.com) and Directory on Server. Specify “YOU” for the directory on the local server. Just like the YOU server, the scheduling of updates on the client can be done manually or at regularly scheduled intervals.

VIII. Alternatives

As with anything in the UNIX world, there are multiple ways to accomplish almost anything. These are just a few of the alternatives to consider in Linux installation on the JS20’s.

VNC or SSH

See *Section I, Remote Access to the JS20*.

Power control of the JS20

When using CSM, the rpower command is the most convenient command for managing the power control of the JS20. However, the following power control alternatives are also available.

In the other install scenarios, we used the MM web interface to power on/off/restart the blades. There are also “power” commands available in the SOL session, via the MM CLI. See the following command sequence:

```
system> power -h
power -off|-cycle {-c}|-state {-post}|-on {-c}
Controls target power for blades and switch modules
-off: power off
-cycle: power off, then on
-state: display power state
-post: details the health status of the module (used on switches)
```



```
-on: power on
-c: enter console mode at power on (used on blades with -on or -cycle)
system> power -state -T blade[1]
On
system> power -off -T blade[1]
OK
system> power -state -T blade[1]
Off
system> power -on -T blade[1]
OK
system> power -state -T blade[1]
On
```

IX. 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/BOOTP problems

The install server (or CSM MS) should have DHCPDISCOVER, DHCPPOFFER, DHCPREQUEST, DHCPACK log entries for each install target in /var/log/messages.

1. Make sure /etc/sysconfig/dhcp contains the line: DHCPD_BINARY="/usr/sbin/dhcpd.lpf"
2. Make sure there is only one DHCP/BOOTP server on the subnet.

TFTP problems

If you encounter problems TFTP'ing the install kernel from the install server, it may be that the TFTP daemon on the install server is not capable of handling a number of parallel sessions. In this case, use either the atftp or tftp-hpa TFTP packages.



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 `/ftpbboot`.

CSM

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

```
probemgr -p ibm.csm.ms -l 0    (lower case L)
```

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 current external list:
<http://techsupport.services.ibm.com/server/csm/documentation/listinfo.html>
- the current internal list and recent archives: <https://w3.opensource.ibm.com/projects/csm/>
- the old internal list and its earlier archives:
<http://lxsvr4.ppd.pok.ibm.com/cgi-bin/lwgate/CSMDEV/>



X. References

BladeCenter JS20 documents:

1. BladeCenter JS20 Type 8842 Installation and User's Guide,
<http://www.ibm.com/pc/support/site.wss/document.do?Indocid=MIGR-54526>
2. IBM eServer BladeCenter - Management Module User's Guide,
<http://www.ibm.com/pc/support/site.wss/document.do?Indocid=MIGR-45153>
3. IBM eServer BladeCenter – Planning and Installation Guide,
<http://www.ibm.com/pc/support/site.wss/document.do?Indocid=MIGR-53670>
4. Serial over LAN Setup Guide – IBM eServer BladeCenter and BladeCenter T,
<http://www.ibm.com/pc/support/site.wss/document.do?Indocid=MIGR-54666>

CSM documents:

5. CSM for Linux 1.4 Administration Guide,
<http://publib.boulder.ibm.com/clresctr/docs/csm/linux/200408/am7al103/am7al10302.html>
6. CSM for Linux 1.4 Planning and Installation Guide,
<http://publib.boulder.ibm.com/clresctr/docs/csm/linux/200408/am7il103/am7il10302.html>
7. CSM for Linux 1.4 Command and Technical Reference,
<http://publib.boulder.ibm.com/clresctr/docs/csm/linux/200408/am7cl102/am7cl10202.html>

SUSE LINUX documents:

8. "AutoYaST Automatic Linux Installation and Configuration with YaST2", Anas Nashif
<http://www.SUSE.com/~nashif/autoinstall/9.0/html/index.html>
9. Installing SUSE LINUX SLES8 SP3 on a BladeCenter JS20 whitepaper,
ftp://ftp.software.ibm.com/pc/pccbbs/pc_servers_pdf/bctrjs20sles8install090704.pdf
10. Installing SUSE LINUX Enterprise Server 9 on an IBM eServer BladeCenter JS20 whitepaper,
ftp://ftp.software.ibm.com/pc/pccbbs/pc_servers_pdf/bctijs20sles9installwp_101204.pdf



Page 43 of 43

Installing SUSE LINUX Enterprise
Server 9 SP1 on an
IBM @server
BladeCenter JS20



©IBM Corporation

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

Produced in the United States of America
August 2005
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, and BladeCenter 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>.

UNIX is a registered trademark of The Open Group in the United States, other countries or both.

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

Windows is a registered trademark of Microsoft Corporation in the United States, other countries, or both.

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>