



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

October 12, 2004

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

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

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

Aaron T. Bolding  
IBM Corporation  
atboldin@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 .....	4
Using a Virtual Network Computing (VNC) session .....	5
II. Manual installation from CD .....	6
Setup to install from CD .....	6
Start the installation .....	6
Install using the YaST screens .....	7
III. Manual installation over network .....	8
Setup of the install server .....	8
DHCP Configuration .....	8
TFTP Configuration .....	10
NFS Configuration .....	10
Building the installation tree .....	10
Setup to install over network .....	12
Start the installation .....	12
Define installation settings using linuxrc .....	13
Install using the YaST screens .....	13
IV. Auto-installation over network via OpenFirmware .....	15
Setup of the install server .....	15
Create the AutoYaST control file .....	15
Install using OpenFirmware .....	17
V. Auto-installation over network, patch install kernel .....	18
Setup of the install server .....	18
Patch kernel options on the install server .....	19
Setup to install over network .....	19
Install using patched kernel .....	19
VI. Auto-installation over network using CSM .....	20
CSM management server (MS) setup .....	21
CSM software package from website .....	21
CSM software package from CDs .....	22
Accept the product license .....	22
Define the Nodes to the MS .....	24
CSM SLES 9 installation setup .....	28
Install the OS to the nodes .....	29
Monitoring with CSM - optional .....	31
VII. YaST Online Update .....	33
YOU Server setup .....	33
YOU Client setup .....	33
VIII. Troubleshooting .....	33
Serial Over LAN problems .....	33
DHCP problems .....	34
TFTP problems .....	34
CSM .....	34
IX. References .....	35



## I. Introduction

This whitepaper explains how to install SUSE LINUX® Enterprise Server 9 onto an IBM eServer® BladeCenter™ JS20 blade 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 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 more systems with varying configurations, or short-lived 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**

This document assumes the Serial Over LAN (SOL) 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 (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.



## ***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:

- eth0 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 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>). 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***

This section describes two mechanisms for accessing the JS20, one for text-based installation and another for GUI-based installation.

### **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/PASSWORD (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 YaST install screens. The free telnet client PUTTY seems to work much better. PUTTY is available on CD #1, in the /dosutils/putty directory. PUTTY can also 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 YaST install 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.

## Using a Virtual Network Computing (VNC) session

To control a remote X-server and therefore enable the use of a GUI for the SLES 9 installation, setup a VNC client system and VNC server to run on the target install system.

### Installing the VNC client

To install a VNC client on a Windows system, copy the tightvnc\*setup.exe file from the /dosutils/tightvnc directory of the SLES 9 CD #1. Run tightvnc\*setup.exe to install the TightVNC client package. Note, to start the VNC client, click on the “TightVNC Viewer” icon (vncviewer.exe).

To install a VNC client on a SLES (ie. UNIX®) system, download the vnc RPM package. Install the package using the YaST2 “Install/Remove Software” option. Note, the VNC client program is vncviewer.

### Starting the VNC client

Later, during the SLES 9 install process, YaST2 will prompt you start the VNC client in order to connect to the VNC server. YaST2 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 UNIX, run the program vncviewer.



Enter the VNC password when prompted. Next, you should see the YaST screen prompting you to select a Language.

### Launching the VNC server on the install target

To setup VNC server to run on the target install system, options must be passed to the kernel. This is accomplished when yaboot presents the “boot:” prompt. At the “boot:” prompt, enter “install vnc=1 vncpassword=PASSWORD” (just a sample password).

## II. Manual installation from CD

This section describes the process of installation from the SLES 9 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 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 SLES 9 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”.

In this example, we'll kick off the install via an SOL session, then use the VNC client GUI capability to more easily maneuver through the YaST2 screens. Then we'll return to the SOL session after YaST2 has completed.

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.



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, 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 the Introduction, Starting the VNC Client) in order to start the GUI-based installation.

## ***Install using the YaST screens***

Next, you should see the YaST screen prompting you 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, timezone 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”. When prompted, insert the correct SLES 9 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, 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 the Introduction, Starting the VNC Client) in order to start the GUI-based installation.

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 your: Network, Online Update, Services and Users. After which 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 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.

SLES 9 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 "shutdown -r now". The server should reboot successfully from the harddrive, 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 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, TFTP, NFS and the install tree satisfy the requirements for a network-based install server. In verifying these SLES 9 install procedures, we actually reused an install server with SLES 8 SP3 installed. Therefore, the instructions for setting up the install server have a SLES appearance.

#### **DHCP Configuration**

First, the install server needs to be setup as a BOOTP/DHCP server. The DHCP server package that comes with SLES 8 or SLES 9 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 tftp line in /etc/inetd.conf.
  - 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 start the DHCP daemon (dhcpd), run:

```
# rcdhcpd start
```

For host name 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, start YaST by running “yast2” from the command line. Then in “Network/Basic” select “Start/stop services (inetd)”. Enable inetd and select Next. Select the “tftp” service and click on “Activate or Deactivate”. Verify that status column for the tftp service indicates “Active”, then click on Finish.

Then we need to copy the `/install` file from the SLES 9 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 `tftp` line in `/etc/inetd.conf`. 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/install
```

## NFS Configuration

Next, we need to setup the install server as an NFS server. To do this, start YaST by running “yast2” from the command line. Then in “Network/Advanced” select “NFS server”. Select “Start NFS server” and select Next. Under “Directories to export..”, click on “Add directory”, specify “/install-tree” and click OK. Let YaST2 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 YaST2 for Hosts wildcard and Options. Then click on Finish and confirm with a “yes”.

Note, you can use the “`rcnfs stop|start|status|restart`” command to further control the NFS daemon.

## 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/sles9 directory. Note, the /install-tree/sles9 directory is used to accommodate other distributions (ie. SLES 8 in the /install-tree/sles8). If you're only installing SLES 9, 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 SLES 9 CDs to disk. Be sure to specify the "periods" in the commands that follow.

```
# mkdir -p /install-tree/sles9/sles/CD1 ; cd /install-tree/sles9/sles/CD1
# umount /media/cdrom
```

insert SLES 9 CD #1

```
# mount /dev/cdrom /media/cdrom ; cp -a /media/cdrom/* . ; umount /media/cdrom
# mkdir -p /install-tree/sles9/core/CD1 ; cd /install-tree/sles9/core/CD1
```

insert SLES 9 CD #2

```
# mount /dev/cdrom /media/cdrom ; cp -a /media/cdrom/* . ; umount /media/cdrom
# mkdir -p /install-tree/sles9/core/CD2 ; cd /install-tree/sles9/core/CD2
```

insert SLES 9 CD #3

```
# mount /dev/cdrom /media/cdrom ; cp -a /media/cdrom/* . ; umount /media/cdrom
# mkdir -p /install-tree/sles9/core/CD3 ; cd /install-tree/sles9/core/CD3
```

insert SLES 9 CD #4

```
# mount /dev/cdrom /media/cdrom ; cp -a /media/cdrom/* . ; umount /media/cdrom
# mkdir -p /install-tree/sles9/core/CD4 ; cd /install-tree/sles9/core/CD4
```

insert SLES 9 CD #5

```
# mount /dev/cdrom /media/cdrom ; cp -a /media/cdrom/* . ; umount /media/cdrom
# mkdir -p /install-tree/sles9/core/CD5 ; cd /install-tree/sles9/core/CD5
```

insert SLES 9 CD #6

```
# 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/sles9
# 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 .
# 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 with an installation source. The instorder file indicates the order in which to install the installation sources

```
# printf "sles/CD1\t\sles/CD1\n" > yast/order
# printf "core/CD1\t\core/CD1\n">> yast/order
# printf "sles/CD1\t\sles/CD1\n" > yast/instorder
# printf "core/CD1\t\core/CD1\n">> yast/instorder
```

That completes the installation tree:

```
/install-tree/
  sles9/
    sles/
    core/
    control.xml
    content
    boot/
    media.1/
    yast/
      order
      instorder
    yastcfg/
```

## 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.

## 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".

In this example, we'll kick off the install via an SOL session, then use the VNC client GUI capability to more easily maneuver through the YaST2 screens.

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. 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.

## ***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/sles9), 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 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 the Introduction, Starting the VNC Client) 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 the Introduction, Starting the VNC Client) in order to start the GUI-based installation.

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 your: Network, Online Update, Services and Users, etc... Setup your network configuration at this time. After which 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 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.

SLES 9 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 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, TFTP, NFS and the install tree just 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 YaST2 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 SLES 8 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 “yast2” from the command line. Then in “Misc” select “Autoinstallation” (or run “yast2 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 SLES 9 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 yast2.

Note, when generating an AutoYaST control file from a SLES 8 system, to be used to install SLES 9 servers, the <bootloader> section must be removed from the SLES 8-based control file.

A sample JS20 AutoYaST control file is provided for download at <http://w?.ibm.com/sales/systems/portal/some???where/sample-JS20-autoyast.xml> Save this file to the /install-tree/sles/yastcfg directory, 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.

Note for the auto-installation procedures we won't start a VNC session, since there's nothing interactive occurring.

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

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

We need to issue the following "boot" command:

```
0 > boot /pci@8000000f8000000/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/sles9/yastcfg/ install=nfs://192.168.70.50/install-tree/sles9
netdevice=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 2<sup>nd</sup> 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 2<sup>nd</sup> 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
- arguments passed to the install kernel are:



- o location of AutoYaST control file = "autoyast=nfs://192.168.70.50/install-tree/sles9/yastcfg"
- o location of install tree = "install=nfs://192.168.70.50/install-tree/sles9"
- o 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 "yast2" 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.

SLES 9 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 "shutdown -r now". The server should reboot successfully from the harddrive, 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 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, 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 SLES 9 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 SLES 9 called `mkzimage_cmdline`. This utility is located in the `/ppc/netboot` directory of SLES 9 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/sles9/yastcfg/  
install=nfs://192.168.70.50/install-tree/sles9 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 the Introduction). 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

Note for the auto-installation procedures we won't start a VNC session, since there's nothing interactive occurring.



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 I, 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. If the SOL session is started soon enough, BOOTP requests will interrupt the SOL session and the session will terminate. 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.

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 “yast2” 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.

SLES 9 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 “shutdown –r now”. The server should reboot successfully from the harddrive, with network connectivity. The installation is complete.

## VI. Auto-installation over network using CSM

Note, at the time of the writing of this whitepaper, CSM 1.4.0.1 does not officially support SUSE LINUX SLES 9. SLES 9 support is targeted for CSM 1.4.0.3.

This section describes how to populate the blade servers using Cluster Systems Management (CSM) for Linux V1.4.0.1. This install option 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 SLES 9 (AutoYaST) 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
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

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/>.

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-1.4.0.1.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 SLES 9 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!".

## 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 SLES 9 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 = 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. use the lshwinfo command to gather attributes from the HW control point
2. transform the lshwinfo output to a node definition file format
3. make alterations to the node definition file
4. issue the definenode 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:

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

The resulting hwinfo1 file should look like this:

```
# Hostname::PowerMethod::HWControlPoint::HWControlNodeId::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 column is separated 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::HWControlNodeId::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= \  
InstallDistributionVersion=9 > nodedef1
```

Note, the “InstallDistributionVersion=9” attr=value pair is only necessary if the CSM MS is not running the same distribution version as the installation target. In our scenario, the CSM MS is a SLES 8 system. It may also be necessary to provide “InstallDistributionName”, if the MS and installation target are not the same distribution.

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=SLES  
InstallDistributionVersion=9  
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=SLES  
InstallDistributionVersion=9
```



```
InstallOSName=Linux  
InstallPkgArchitecture=ppc64  
ManagementServer=ms.bclab.com  
PowerMethod=blade
```

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
```

... 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 = SLES  
InstallDistributionVersion = 9  
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 = SLES  
InstallDistributionVersion = 9  
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 SLES 9 installation setup**

Now we need to supply the SLES 9 GA CDs to CSM so it can build the install tree and the AutoYaST control file.

A couple other things first though. Unlike other installs 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.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 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 SLES 9 CDs to CSM, run the following command:

```
# copycds InstallDistributionName=SLES InstallDistributionVersion=9
```

and insert the next SLES 9 CD when prompted. As CSM copies the contents of the CD, the % complete should be displayed in the lower left of the window. Note, if the AutoYaST control file template was not modified, CSM will just ask for SLES 9 CDs #1, #2 and #3.



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.com	PreManaged	Not Installed
blade2.bclab.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. The “-P” option indicates to create the control file for each node with Mode=PreManaged. The output of the csmsetupyast command can be found in /var/log/csm/csmsetupyast.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 AutoYaST configuration files in /csminstall/csm/1.4.0/Autoyast.SLES9

## ***Install the OS to the nodes***

The next step is to install SLES 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 = SLES
DistributionVersion = 9
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 = SLES
DistributionVersion = 9
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 are 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"	"rconsoleUpdateResponse"	"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 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

The 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.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 “yast2” from the command line. Highlight “Software”, 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 regulary scheduled basis.

### ***YOU Client setup***

On the YOU client side, start YaST by running “yast2” 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.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. Troubleshooting

### ***Serial Over LAN problems***

If Serial Over LAN sessions aren’t showing you the expected output, here’s 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 that 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***

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/>



## IX. References

1. BladeCenter JS20 Type 8842 Installation and User's Guide,  
<http://www-307.ibm.com/pc/support/site.wss/document.do?Indocid=MIGR-54526>
2. IBM eServer BladeCenter - Management Module User's Guide,  
<http://www-307.ibm.com/pc/support/site.wss/document.do?Indocid=MIGR-45153>
3. IBM eServer BladeCenter – Planning and Installation Guide,  
<http://www-307.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-307.ibm.com/pc/support/site.wss/document.do?Indocid=MIGR-54666>
5. “AutoYaST Automatic Linux Installation and Configuration with YaST2”, Anas Nashif  
<http://www.SUSE.com/~nashif/autoinstall/9.0/html/index.html>
6. CSM for Linux 1.4 Administration Guide,  
<http://publib.boulder.ibm.com/clresctr/docs/csm/linux/200408/am7al103/am7al10302.html>
7. CSM for Linux 1.4 Planning and Installation Guide,  
<http://publib.boulder.ibm.com/clresctr/docs/csm/linux/200408/am7il103/am7il10302.html>
8. CSM for Linux 1.4 Command and Technical Reference,  
<http://publib.boulder.ibm.com/clresctr/docs/csm/linux/200408/am7cl102/am7cl10202.html>



Page 36 of 36

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

© 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 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>.

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

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

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

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.

IBM, the IBM logo, @server, eServer, BladeCenter 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>.

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

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

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

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