



Using the @server Cluster 1350 Linux Cluster Installation Tool version 3.2

The IBM® @server® Cluster 1350 Linux Cluster Installation Tool (LCIT) version 3.2 provides a graphical interface that presents imported or discovered cluster data information. LCIT 3.2 builds a database representing the hardware resources in the cluster. After the database is created, you can then use LCIT, to assign or change device IP addresses, such as RSA, ISMP, and IPMI addresses, perform diagnostics, update the BIOS code and other firmware, and write configuration files required by XCAT and CSM.

Preinstallation requirements

To install the hardware components into the racks before running the LCIT installation diagnostics, complete the following steps:

1. Populate the rack cabinets with your cluster components.
2. Complete all wiring of the rack components, making sure that any Remote Supervisor Adapter (RSA) devices are in the management VLAN switch ports or on a dedicated switch.
3. Configure all the switches, RSA devices, terminal servers or port servers, and FASTT components, using the applicable documentation.

Running Linux cluster installation tool diagnostics

The installation materials include a LCIT startable CD. With your 1350 Cluster, you also receive tab files that detail the rack hardware configuration. To make sure that you install the cluster components correctly, run the installation diagnostics.

Note: You can find the latest updates for LCIT at <http://publib.boulder.ibm.com/cluster/index.htm>.

To run the Linux cluster installation tool, complete the following steps:

1. In the rear of the rack, unplug all compute nodes.
2. Turn on the management node and insert the LCIT CD into the CD-ROM drive. Workspace 1 opens as a gray screen.
3. Anywhere in the open Workspace 1 window, right-click and select **LCIT 3.2**. The LCIT window opens.
4. Right-click on the IP address **172.20.0.1**, and then select **Add Search Range/Alias** to set the Baseboard Management Controller (BMC). The factory default is 172.29.0.1 with 16-bit subnet.
5. Click **OK**.
6. Click **Configuration** and click **Apply** to write the dhcpd.conf file.
7. Plug in the power cords for all the compute nodes in chronological order. Node numbering increases from the bottom of the rack upward and from left to right in the BladeCenter™ unit.
8. Start each server and make sure that the servers display the management-node Workspace 1 gray screen. This step validates that each server has loaded the ram-disk image from the network.

Note: If you do not see a Workspace 1 gray screen on a server, make sure that the server is configured to boot from the network.

Installing an operating system on the management node

Before an IBM support representative can install LCIT, you must install the operating system on the management node.

To install an operating system, complete the following steps:

1. Turn on the management-cluster node.
2. Insert the Linux software CD-ROM Disk 1 into the CD-ROM drive and reboot the server. Disk 1 is programmed to start automatically and auto-run. The installation program window opens.
3. During this installation you are given the option to partition your hard disks. Configure the LCIT partition for 200 MB and set up the partition on SCSI drive array 5 (sda5).

Configuring remote connections

To configure the management node to accept remote connections, you must modify the syslog configuration file. To modify the syslog configuration file, complete the following steps:

1. Open a console window and type: **vi /etc/sysconfig/syslog**.
The syslog configuration file opens.
2. Change the line `SYSLOGD_OPTIONS= m 0` to:
SYSLOGD_OPTIONS= -m 0 -r
3. Restart the syslog daemon to save the changes, by typing:
service syslog restart
4. Mount the LCIT CD using one of the following commands:
 - For Red Hat Linux, type: **mount /mnt/cdrom**
 - For SLES, type: **mount /media/cdrom**
5. Copy LCIT to the root directory, by typing: **cp LCIT /root**.
6. Install RPM, by typing: **rpm -uvh lcit-* .rpm**.

Installing Linux binaries

You must download the Linux PXE boot binaries, untar them, and copy the pxelinux.0 and memdisk images to the /TFTP boot directory. The Linux binaries can be downloaded from <http://www.kernel.org/pub/linux/utils/boot/syslinux/>.

Network default addresses

Use the following tables to reference the LCIT network device defaults.

Table 1. Network address defaults

IP address	Host name	Component
172.20.0.1	mgt.cluster.com	Management node eth0 (cluster VLAN)
172.30.0.1	mgt1.cluster.com	Management node eth1 (management VLAN)
172.20.1.1	storage001	Storage nodes
172.30.2.1	triton001	FAST storage controllers
172.20.101.1	node001	Compute nodes
172.20.4.1	user001	User nodes
172.20.10.1	myri001	Myrinet switch
172.30.20.1	ts001	Terminal servers
172.30.30.1	rsa001	Remote System Access devices

Table 1. Network address defaults (continued)

IP address	Host name	Component
172.29.101.1	BMC001	Baseboard management controller (IPMI devices)
172.30.40.1	cisco3550-001	10/100 Cisco switches
172.30.50.1	cisco3508-001	Gigabit Cisco switches
172.30.60.1	apc001	APC
172.30.70.1	rcm001	RCM
172.30.80.1	cisco4003-001	4000 Cisco series switch
172.30.101.1	sm001/mm001	BladeCenter switch modules and management modules

Table 2. Compute node IP address defaults

Rack	Network	Rack number	Node
Rack 1	172.20	101	1 to 84
Rack 2	172.20	102	1 to 84

Table 3. Baseboard management controller IP address defaults

Rack	Network	Rack number	Node
Rack 1	172.29	101	1 to 40
Rack 2	172.29	102	1 to 40

Note: Node numbering increases from the bottom of the rack upward and from left to right in the BladeCenter units.

Table 4. Manufacturing user IDs and password defaults

Hardware	User ID	Password	Comments
APC switch	apc	apc	
ITouch terminal	access	system	
Cisco switch	press Enter for the user ID	ibm 1350	
FASTT	press Enter for the user ID	infiniti	
RSA	USERID	PASSWORD	The 0 in PASSWORD is a zero.

Table 5. Operating system user IDs and password defaults

Operating system	User ID	Password
Management server	root	ibm1350
Nodes	admin	cluster
Admin	root	cluster
Root	cluster	cluster

Note: To watch the system log, open another console window and type: `tail -f /var/log/messages`.

Starting DHCP, NFS, and xinetd for the first time

The first time you start LCIT, you need to start several of the network services. To start these network services, complete the following steps:

1. To start LCIT, at the root console prompt, type: `/LCIT/bin/lcit`. The LCIT program window opens.
2. Click **Configuration** and review the LCIT device configuration information that is listed in each tab.
3. If the default configuration information is applicable to your network, click **Apply**. This replaces the LCIT network configuration `/etc/xinetd.d/tftp` file with the new configuration information.
4. At the root console prompt, type: `service dhcpd restart`. This command restarts the DHCP service.
5. At the root console prompt, type: `service nfs restart`. This command restarts the NFS service.
6. At the root console prompt, type: `service xinetd restart`. This command restarts the xinetd service.

Configuring LCIT services to autostart

After you configure LCIT services for the first time, you can configure LCIT services to automatically start upon rebooting the server. To use LCIT, you must log into the root directory.

Perform the following steps to configure LCIT services to automatically start upon rebooting the server:

1. At the root console prompt, type: `setup`.
2. Using the keyboard arrow keys, scroll down to the system services section, and press the space bar to select **system services**.
3. Scroll down and use the space bar to select the following services:
`dhcpd`
`nfs`
`xinetd`
4. Using the Tab key, select the **OK** button, and press **Enter**.
5. To verify that the services have started, at the root console prompt, type:
`service dhcpd restart`
`service nfs restart`
`service xinetd restart`
6. To start LCIT, at the root console prompt, type: `/LCIT/bin/lcit`. The LCIT window opens.

Discovering cluster devices

To continue with installing LCIT you must discover the Cluster 1350 devices. To discover cluster devices, complete the following steps:

1. From the **View** menu, select **Discover**. The discovery process polls for all Ethernet controllers and lists the discovered Ethernet controllers in the Discovery/Controller tab.
2. Click **Next** to continue the discovery process. The discovery wizard polls all configured search ranges for cluster device IP and MAC addresses and populates the IP Addresses tab and Rack Configurator tab.

Note: You can control the speed at which LCIT polls for IP device information by changing the address request protocol (ARP) Group number. If the discovery process is not reporting information, decrease the ARP Group number. The default ARP Group number is 6.

3. Click **Next** to discover all associated IPMI devices. A list of IPMI devices opens in the IPMI tab.
4. Click **Next** to discover all associated RSA devices. A list of RSA devices opens in the RSA tab. When the process is complete, the Next button changes to Scan.

Note: After the first discovery, you can run a scan again to poll for new devices that may have been added to the cluster or are not working correctly.

5. After the discovery process is complete, from the **View** menu, select **Cluster Hardware Survey** to review the inventory of devices found during the discovery process.

Note: Devices that are listed in red have duplicate IP addresses. Devices that are listed in blue are Dynamic Host Control Protocol (DHCP) devices.

6. To review cluster statistical information:
 - In the Cluster Tree panel, single-click a device. Statistical information about the device displays in the right-panel of the interface. This gives access to information about individual BIOS code levels, flash levels, and many other statistics. Ethernet Controllers display dynamic information.
 - To review the rack configurator image and associated information, in the Cluster Tree panel, double-click a device. The rack configurator image displays in the right-panel of the interface. The rack configurator image displays the rack and its devices in their physical order and placement within the rack cabinet.

Resetting remote supervisor adapter nodes

To reset Remote Supervisor Adapter (RSA) cluster nodes, perform the following steps:

1. From the **View** menu, select **Cluster Hardware Survey**.
2. In the **Cluster Tree** panel, click the **Ethernet Controller** expand icon to list the RSA nodes. A list of devices opens.
3. Click the **RSA** node that you want to reset and perform one of the following actions:
 - To turn on the power to all devices in the RSA chain, click **Power On**.
 - To turn off the power to all devices in the RSA chain, click **Power Off**.
 - To reset all ISMP processors in the RSA chain, click **Reset Remote ISMPs**. This turns off the power for all ISMP devices in the RSA chain then turns the power on.
 - To reset and reboot an RSA device, click **Reset RSA Device**.
 - To change an RSA IP address, if DHCP is disabled, click the up or down arrows for the portion of the IP address that you want to change. The up-arrow increases the number and the down-arrow decreases the number. If DHCP is enabled, this function is not accessible.

Note: After the RSA nodes are located, you can verify the nodes are operating properly by hovering your mouse over the RSA item in the LCIT Cluster Tree panel and visually inspecting the RSA module in the rack for a lit blue LED. If the blue LED is lit, the device is online and communicating with the cluster. If the blue LED is not lit, there is a possible problem with that device. This is also helpful for visually locating a device in a rack.

Resetting remote supervisor adapter service processors

The Remote Supervisor Adapters (RSAs) in the @server 325 server supports Intelligent Service Management Processors (ISMPs). To reset the host server associated with the ISMPs, complete the following steps:

1. From the **View** menu, select **Cluster Hardware Survey**.
2. In the **Cluster Tree** panel, click the **Ethernet Controller** expand icon. A list of devices opens.
3. Click the **RSA** expand icon, if applicable, and single-click the **ISMP** that you want to reset. The ISMP tab opens.
4. Perform one of the following actions:
 - To turn off an ISMP, click **Power**. This is a toggle switch that turns off a device that is turned on and turns on a device that is turned off. You can watch the status text window to observe the result of your action.
 - To turn off and turn on an ISMP, click **Reset**.
 - To change the IP address of an ISMP, click the up or down arrows for the portion of the IP address that you want to change. The up-arrow increases the number and the down-arrow decreases the number. If DHCP is enabled, this function is not accessible. To enable or disable DHCP, click the **DHCP** tab and reset the state.

Note: After the ISMP processors are located, you can verify their nodes are operating properly by hovering your mouse over the ISMP device in the Cluster Tree panel and visually inspecting the module in the rack for a lit blue LED. If the blue LED is lit, the device is online and communicating with the cluster. If the blue LED is not lit, there is a possible problem with that device. This is also helpful for visually locating a device in a rack.

Resetting intelligent platform management interface nodes

To reset the operational state of the cluster Remote Supervisor Adapter (RSA) Intelligent Platform Management Interface (IPMI) nodes, perform the following steps:

1. From the **View** menu, select **Cluster Hardware Survey**.
2. In the **Cluster Tree** panel, click the **Ethernet Controller** expand icon. A list of devices opens.
3. Single-click the **IPMI** device that you want to reset. The IPMI device tab opens.
4. Click the **Power** tab and perform one of the following actions:
 - To turn off an IPMI device, click **Power Down**. This is a toggle switch that turns off a device that is turned on and turns on a device that is turned off. You can watch the status text window to observe the result of your action. When power is turned on, the Power button is pushed inwards.
 - To turn on an IPMI device, click **Power Up** to turn on the selected IPMI device.

Note: After the IPMI nodes are located, you can verify the nodes are operating properly by hovering your mouse over the IPMI device in the Cluster Tree panel and visually inspecting the module in the rack for a lit blue LED. If the blue LED is lit, the device is online and communicating with the cluster. If the blue LED is not lit, there is a possible problem with that device. This is also helpful for visually locating a device in a rack.

Assigning a node to a rack position


Using LCIT, you can assign a node to a specific position in a rack. By default, LCIT creates an unpopulated Rack Configurator template. After the initial LCIT discovery, device type, and address information is collected and is listed in the Rack Configurator tab.

To assign a node to a rack position, complete the following steps:

1. Click the **View** menu and select **Discover**; then click the **Rack Configurator** tab.
2. In the **Cluster Tree** panel, double-click and hold a device while dragging the device into the Rack Configurator image into the desired physical rack location. Repeat this step until all rack devices are listed in the Rack Configurator.
3. Repeat steps 1 and 2 to populate each new rack template that you want to create.
4. After you complete steps 1 through 4, you must reset the management module for each rack and each RSA device to complete the IP address reassignment.
 - a. To reset a management node, use a straightened paper clip to press the recessed button on the management module in the rear of the chassis. Repeat this step for each management module in each rack.
 - b. To reset an RSA device, in the **Cluster Tree** panel, single-click an **RSA** device and in the right-panel click **Restart**.

Note: To restart standard nodes, you can disconnect and reconnect the power cords, or in the LCIT interface, in the Cluster Tree panel, single-click the device, and click **Power Down**, then click **Power Up** to restart the node.

Figure 1. Management module IP reset button



Adding an Ethernet controller device

Using LCIT you can add Ethernet controller devices to the Cluster 1350. This is helpful for planning and configuring a device before it is physically installed in the rack.

To add an Ethernet controller device, in the **Cluster Tree** panel, right-click an Ethernet controller; then select the device to be added. The new device is listed in the Cluster Tree panel and in the device information set. A single-click on an Ethernet controller device displays the dynamic device information for that controller.

Creating a group

Using LCIT you can create groups of like devices, groups of IP address ranges; groups that aid in focusing on specific cluster management. You can also create groups within groups without restriction. For example, you can group RSA devices and group IPMI devices into a subgroup.

Note: Devices can not be moved across groups and can not belong to more than one group at a time.

To create a group, complete the following steps:

1. In the **Cluster Tree** panel, right-click a device, and select **Group**.
2. In the **Cluster Tree** panel, right-click the **Name** column, type the name of the group, and press **Enter**. The new group is listed in the Cluster Tree panel.

Adding a device to a group

To add a device to an existing group, in the Cluster Tree panel, double-click, hold a device, and drag it into the new group. The device is listed under the group.

Note: You can add new devices to the group. Right-click the existing group and select the device to be added. Devices can not be moved across groups and can not belong to more than one group at a time.

Related publications

For more information about the Cluster 1350, see the following Web sites:

- <http://publib.boulder.ibm.com/cluster/index.htm>
- <http://techsupport.services.ibm.com/linux/techKnow>

Fifth Edition (June 2004)

Printed in the U.S.A.

IBM, @server, ServeRAID, and the e-business logo are trademarks of the IBM Corporation in the United States, other countries, or both.

© Copyright International Business Machines Corporation 2004. All rights reserved.

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

(1P) P/N: 25K8421

