

SGI™ Origin™ 3000 Series L3 Controller 1.3 Update Guide

Overview

The L3 software consists of four components:

- Kernel software
- L3 driver
- L1/L2 firmware utilities
- L3 install tool

This product provides a system control platform for field service and system maintenance applications. The features include:

- Advanced diagnostic support
- Maintenance of the SGI Origin 3000 system serial number
- Maintenance of the SGI Origin 3000 router port enable
- Console for L1 and L2 controllers

Related Information

For more information on the SGI L1, L2, and L3 controllers, refer to the following documents:

- *SGI Origin 3000 L3 Controller Installation Guide*, part number 007-4363-001
- *SGI Origin 3000 Series Owner's Guide*, part number 007-4240-001

Installing the Diagnostic Software

To install the diagnostic software, perform the following steps:

1. Log onto the L3 platform as root.
2. Insert the Origin 3000 Series L3 Controller 1.3 CD-ROM into the CD-ROM drive.
3. Enter the following commands to load the diagnostic software:

```
# cd /mnt/cdrom/RPMS  
# 13_install
```

4. Remove the CD-ROM from the CD-ROM drive.
5. Reboot the system:

```
# reboot
```

Software Installation on Systems with Preinstalled Red Hat™ Linux®

In order to install the L3 Controller software on a Silicon Graphics 230 visual workstation that has the Red Hat Linux 6.2 operating system preinstalled during manufacturing, some additional steps must be performed. Instructions for performing these steps can be found in the *SGI Origin 3000 L3 Controller Installation Guide* (publication number 007-4363-001).

Refer to the following sections in the *SGI Origin 3000 L3 Controller Installation Guide* for detailed installation instructions:

- Installation Overview (p. 2)
- Software Installation Requirements (pp. 2 – 4)
- Installing the Auxiliary Linux Software (p. 18)
- Installing the L3 Controller Software (p. 19)
- Installing the Software from the CD-ROM (pp. 19 – 20)

Note: To configure the network (using the Network Configurator GUI) after the workstation has booted to OS, refer to the “Configuring Network Settings” section (pp. 12 – 14). To configure user accounts (using the User Configurator GUI), refer to the “Configuring the Account”

section (p. 15). To connect the Silicon Graphics 230 visual workstation to an SGI Origin 3000 series server, refer to the sections that are appropriate for your system configuration (pp. 21 – 27).

New Features and Changes

The following new features and changes are included in this release.

L1 Firmware

Note: Upgrading to this version of L1 firmware (version 1.4.1) will set the voltage margins to `DEFAULT` on all C-bricks. If you are running with special voltage margins on any C-bricks (margins other than `NORMAL`), record the current settings *before* you upgrade the firmware. Then, after you upgrade, set the voltage margins to their desired settings by using the `L1 margin` command. If no special settings were previously in use, use the `L1 margin normal` command to set the voltage margins back to the `NORMAL` setting.

Added Commands

<code>cfg v</code>	Added as an alias for <code>cfg verbose</code> .
<code>link</code>	Examines the state of all NI and II ports on the brick, and reports an error for each unexpected state (which indicate that the port does not have a link when it should).
<code>margin VTERM</code>	Margins the XIO card termination voltage on X-brick.
<code>pbay env</code>	Reports the state of power bay monitoring on the L1.
<code>pbay env on off</code>	Enables or disables power bay monitoring from the L1.
<code>pbay init</code>	Manually initializes power bay (ICMB) interface.
<code>router spare</code>	Command to correctly configure R-bricks during manufacturing (sets the router port enabling to the correct state).
<code>usb</code>	Displays debugging statistics about the L1-to-L2 connection.

Changed Commands

<code>cfg</code>	Now reports bricks connected but not configured (which usually indicates a connection error).
<code>pci</code>	Added PCI hot-swap restrictions for 33MHz cards in a 66MHz bus.
<code>sel s a b c d</code>	Correct subchannel number (0 – 3) is now set. (Previous versions set an invalid subchannel number.)
<code>serial</code>	Strips off any trailing tab in brick serial number.

Added Features

Added support for 500MHz R14K IP35 PIMMs.

The L1 can now detect PIMM type changes across L1 reboots and reset the necessary parameters appropriately (that is, when upgrading to 500MHz R14K IP35 PIMMs).

Added support for PCI hot swap. Added PCI hot-swap restrictions for 33MHz cards in a 66MHz bus:

- Cannot now add a 33MHz card to an empty 66MHz bus.
- Cannot remove the last card from a 66MHz bus that is running at 33MHz.

- These restrictions are necessary to prevent the XBridge PCI clock from switching from 33MHz to 66MHz in operational mode and hanging the XBridge (which causes IRIX to crash).

UST clock source on an I/O brick will always default to an active port.

Margining of the XIO card termination voltage on the X brick is now supported by the standard `margin` command.

Eliminated loss of console data between the L1 and the L2 that was caused by the L2 not servicing the L1 often enough.

Data passed between L1 and L2 is now prioritized. This should eliminate "timeout" responses from commands at the L2 due to heavy system console traffic.

Added monitoring of power bay status (DPS, AC port, and DC port failures). Each L1 probes its power bay's status every 5 minutes and reports any failures. In this version, this monitoring is disabled by default, and must be enabled with the `pbay env` command. (Note: This functionality is dependent upon having the correct version of the server management card in the power bay).

Changed the way messages are displayed on the front-panel display. There are now three types of messages:

- FAULT/WARNING messages are displayed when some environmental or power condition occurs in the brick. These messages have the highest priority and are displayed until the condition is corrected.
- ACTION messages are messages that are displayed to indicate an action that the user has taken. These messages are displayed for 15 seconds after the action occurs, and then are removed. (For example, when the user resets a brick, the message "Reset done" is displayed for 15 seconds on that L1's front-panel display.
- INFO messages display the state of the machine (such as "Powered Up") and are displayed when no other types of messages are being displayed.

The default message on the front-panel display is now "Powered Up" or "Powered Down", rather than just "L1 running".

Other Changes

Improved the resiliency of L1-to-L2 USB connections by adding the capability to timeout transfers on the L1 side and attempting to reconnect with the L2 when the connection is lost.

Fixed bugs in C-to-C and C-to-I serial communications that could cause connections to be lost until the C-brick or I-brick was power cycled.

Updated the setting of I/O port speeds for various I/O brick part numbers. Part numbers 030-1565-001L, 030-1565-004, 030-1454-005, 030-1557-007, and 030-1710-* all force the I/O port speed to 600Mhz but report the setting as 400MHz.

I/O brick partition numbers are now updated when their attached C-brick partition number changes. (Both the internal value and the value shown on the L1 front-panel display are updated.)

Changed the power-up order of power supplies in an I-brick. Powering up the 5V supply before the 12V supply was contributing to a high I-brick failure rate with certain types of IBM disk drives.

Removed an invalid environmental alert (usually a low fault or low warning) reported during power down of a brick.

Fixed a G-brick environmental monitor issue. (Previous version would power down a G-brick after a temperature warning.)

Fixed system console selection problems on 2-C partitioned configurations.

Fixed a brick serial number mismatch problem (caused by bugs in previous firmware versions).

Bugs Addressed

- 788242 L1 display board not refreshed after NMI, Reset.
- 801806 Need to margin VTERM.
- 806281 ERROR: SerNum:Brick Serial Number mismatch. See log for details.
- 809406 Need power bay report back failures to system controller.
- 809058 New X brick midplane part # requires 600MHz clock.
- 812973 L1 doesn't pass events to bedrock/OS/ESP.
- 817474 UST clock sourcing should default to active port.

L2 Firmware

Added Commands

<code>cfg s summary</code>	Provides a count of L2s and each brick type in the system and reports any connected but unconfigured USB connections or C-C/C-I connections.
<code>env s</code>	Reports any environmental warning or fault conditions on any bricks in the system.
<code>pwr s</code>	Provides a count of bricks powered up and down and a count of bricks at each margin state (high, normal, or low).
<code>sel reset</code>	Clears any console brick/subchannel overrides in this l2term (SMP) session.

Changed Commands

<code>cfg</code>	Output is sorted by rack/slot number. Reports unconfigured USB and RS-422 (C-C and C-I) connections.
<code>loopback</code>	Fixed so that the L2 emulator should no longer crash.

`l2find` Responses are now sorted by rack number and then by IP address.

`sel s a|b|c|d` Correct subchannel number (0 – 3) is now set. (Running this command in previous versions caused an emulator core dump.)

Added Features

Traffic between L2s is now prioritized. This should eliminate “timeout” responses to L2 commands (particularly during times of heavy console traffic, such as powering up).

When entering commands at the L2> prompt, CTRL-C can be used to clear the current command line and to abort the current command (when waiting for a response or displaying the responses).

Modified the power-up sequence to improve speed on large systems:

- All routers, one at a time.
- All I/O bricks in parallel.
- All C-bricks in parallel.

Other Changes

Added support for systems up to 512 processors (16 L2 Controllers) and should scale to even larger systems.

Added support for rack numbers up to 256.

Command-line and L2 command response should no longer become unacceptably slow during times of heavy console traffic (such as powering up).

Connections between L2s should be significantly more stable.

Improved resiliency and error recovery in L2-to-L1 USB connections. (Connection problems between an L1 and its L2 should no longer make the system controllers unacceptably slow.)

System console selection should work properly, including the ability to select subchannels of NMI'd bricks as necessary.

Error messages now display the rack and slot number of the brick where the error originated (where possible). (Old versions often displayed only the L2's IP address in error messages.)

Fixed bug that caused an L2 to occasionally lose the rack number information for other L2s that it is connected to.

Bugs Addressed

- 811082 Console jumps between same numbered bricks in different racks.
- 812932 System power-up of 128P systems causes errors with L2 1.3.0.
- 813064 L1 controller messages should stay in the partition.
- 815000 L2 thinks partitioning exists when none does.
- 816519 PROM infinite loop.

L3 Software

Added Applications

None.

Changed Applications

- `flashsc` Due to changes in the L1 firmware, older `flashsc` binaries cannot be used to update L1 firmware images in this release. (Older versions should display a warning message informing the user to upgrade to a newer version of `flashsc`.)
Added retry capability when an error occurs when communicating with an L1 or L2 being flashed. `flashsc` will now attempt any flash erase or program operation 3 times before giving up.
- `l2` All changes listed above for the L2 firmware also apply to the L2 emulator.
- `l2find` The list of L2s is now sorted first by system name or SSN, then by rack number, and then by IP address. (Local responses are still returned first).

Added Features

Enhancements to the behavior of the standard system controller command-line options (`--sysname`, `--ssn`, `--l2`, `--interface`):

- When using `--sysname` or `--ssn` to connect to a system, the connection will always be made to the lowest-numbered rack available on that system.
- The `--sysname` and `--ssn` arguments can now include a colon and a rack number after the name or serial number. This restricts the application to connect only to the L2 with the specified rack number. For example, `l2term --sysname mysys:3` will connect to the L2 with the sysname `mysys` and rack number set to 3.

Other Changes

None.

Bugs Addressed

- 812690 Error reported by “flashsc --serial” command.
- 813169 Need to install “flashsc” into /usr/sbin on L3.

©2001, Silicon Graphics, Inc. All rights reserved; provided portions may be copyright in third parties, as indicated elsewhere herein. No permission is granted to copy, distribute, or create derivative works from the contents of this electronic documentation in any manner, in whole or in part, without the prior written permission of Silicon Graphics, Inc. Silicon Graphics and IRIX are registered trademarks, and SGI, the SGI logo, and Origin are trademarks of Silicon Graphics, Inc. Linux is a registered trademark of Linus Torvalds. Red Hat is a trademark of Red Hat, Inc.