Enclosure Mapping in LSISAS-2 Controllers

Systems Engineering Note

S11209, Version 1.0
August 2011

DB05-000290-00

Revision History

Version and Date	Description of Changes
Version 1.0, August 2011	Initial release of this document.

Corporate Headquarters
Milpitas, CA
800-372-2447

Email
globalsupport@lsi.com

Website
www.lsi.com

Document Number: DB05-000290-00

Enclosure Mapping in LSISAS-2 Controllers Systems Engineering Note

In SAS-1 and previous LSI® Fusion-MPT™ storage controllers, the controller firmware mapped each discovered target device to a bus/target ID, using parameters provided by the host driver using the IOCInit message. These parameters informed the firmware of the operating system (OS) constraints on bus/target mapping (such as the maximum number of targets per bus). The OS host driver did not need to handle any bus/target mapping functions.

However, in SAS-2 controllers, bus/target mapping is removed from the firmware, and the mapping is the responsibility of any host driver where the OS uses bus/target mapping to identify devices. Currently, all Windows® OS versions and older (2.4 kernel-based) Linux® OS versions require bus/target mapping.

NOTE Linux kernel 2.6 removed bus/target mapping and requires all drivers to map based on first-come phy location in the topology. To fix the persistency issue in the Linux OS, Linux created the udev facility to provide unique naming to the hard disk drives (HDDs) found during discovery. The Linux 2.6 OS is SAS World Wide Name (WWN) aware. The Windows OS and older Linux kernels are not SAS WWN aware.

Host mapping is the process of mapping devices or Integrated RAID (IR) volumes to a bus/target ID that is meaningful to the host OS. Mapping requirements fall into two basic types, SAS address mapping and enclosure/slot mapping.

In SAS address mapping, a bus/target ID is assigned to a particular target and stays with that target, even if the target is physically moved to a different location in the SAS topology.

In enclosure/slot mapping, a bus/target ID is assigned to a particular target and is dependent on the physical location of the device (typically in a slot within an enclosure). The bus/target mapping is based on the physical controller or expander PHY where the device attaches. If a device is removed and a different device is attached in the same location, the new device is assigned the bus/target ID of the previous device.

1 Device Persistency

Device persistence is the ability for the host driver to maintain the bus/target mapping of devices through controller hard resets and power cycles (system reboots). Because of the SAS discovery process, an original set of bus/target mappings can change if the controller is hard reset and device persistency is not used. The change occurs because of additional devices or removed devices in the topology. A hard reset (or power cycle) can result in different device handles that the firmware uses to identify target devices. When device persistency is implemented, information about the SAS address or enclosure/slot mapping is stored in nonvolatile memory (Driver Persistent Mapping Page 0). This information is available to the driver when the controller re-initializes to let the driver match each SAS address or PHY mapping back to the correct bus/target ID.

2 SAS Address Mapping and Enclosure/Slot Mapping

2.1 SAS Address Mapping Mode

Previously, the LSI standard adapter NVDATA settings used SAS address mapping. This mapping is most beneficial when using direct-attach devices and a smaller number of devices (less than100 devices). As currently defined, the LSI Driver Persistent Mapping Page 0 has a 128-entry limit. If SAS address mapping is used and more than 128 target devices exist, persistent mapping is not possible on all of the devices. The SAS address mapping mode is simple to use with regards to the driver, and it is easy to understand. Any particular drive maps to a bus/target ID, and the drive maintains that bus/target ID no matter where the drive is placed in the SAS topology. The drive address is maintained through device reboots.

2.2 Enclosure/Slot Mapping Mode

Use the enclosure/slot mapping mode if the external storage consists of one or more enclosures. If many targets are used in enclosures, enclosure/slot mapping is ideal because the entries in Driver Persistent Mapping Page 0 are enclosure entries, not device entries. You can attach up to 128 enclosures (including the adapter virtual enclosure) maintained persistently. Limits exist, such as the maximum number of devices the firmware can handle, which limits the number of actual, present enclosures that can be mapped.

Enclosure/slot mapping is the only mode where the LSI firmware multipath capability is supported. The multipath capability is two paths to the same enclosure, and each path detects the enclosure with the same Enclosure Logical ID. The PhyBits (a bitmap representation of the PHYs that connect that path) are different. The SAS address mapping mode does not have a similar multipath capability.

3 Check Your Mapping Mode and Persistence

Use the following process to quickly check the mapping mode and persistence settings on your LSI controller.

1. Open the LSIUtil program.
2. Select Option 9.
3. Enter 1 for the Page Type value. A value of 1 indicates IOC pages.
4. Enter 8 for the Page Number value. A value of 8 indicates page 8.
5. To view the NVDATA version, select 0 for NVRAM. To view the current version loaded that firmware is running, choose 1 for CURRENT.
6. Read offset 0xC.
   - Bit 1
     - 1: Enclosure/slot mapping mode.
     - 0: SAS address hashed mode.
   - Bit 0
     - 1: Persistence is disabled.
     - 0: Persistence is enabled.

4 SAS Address Mapping Mode to Enclosure/Slot Mapping Mode

Use the following process to change the mapping mode.

You must clear the Device Persistent Page 0 and change the mapping mode (change IOC Page 8) in DOS, not in the Windows OS. As soon as you start the Windows OS, the information is written to DPM Page 0, and it is formatted for the current mapping mode. Therefore, before you can run the LSIUtil program, some entries are written.

Changing the mapping mode and then clearing the persistency information is not sufficient because the driver could still write to DPM Page 0 before you change the mapping mode and persistency settings. To make sure that you have a consistent state and a successful mode change, perform these actions using DOS with no underlying driver and no accesses to DPM Page 0.

1. Boot your computer system to DOS.
2. Erase the Persistence region by using SAS2Flash.

   `sas2flsh –o –e 3`

   This action removes any device SAS addresses or hashes from the persistent region.
3. Reset any other configuration page settings that are different from the default NVDATA on the adapter.
   a. Are you using the LSIUtil program to change the mapping mode to enclosure/slot mode? Go to step 4.
   b. Are you updating NVDATA and reflashing the firmware to change the mapping mode to enclosure/slot mode? Go to step 5.
4. Use LSIUtil option 9 to change the IOC Page 8 to place the host bus adapter (HBA) in the enclosure/slot mapping mode (use PageType value 1 and PageNumber value 8).
   a. Set the second bit (bit 1) of the flag field (Offset 0x0C) of IOC Page 8.
   b. When prompted by the LSIUtil program, choose to change the NVRAM settings.
   c. Go to step 6.
5. Modify the NVDATA with ForceUpdate=`true` to IOC Page 8 with the change in step 4a.
   a. Concatenate the NVDATA with the general customer availability (GCA) firmware.
   b. Download the firmware to your HBA.
6. Boot to your OS (Windows, Linux, and so on).

The size of the driver mapping table also must be appropriate. To see all of the devices in all of the enclosures, the mapping table must be large enough to hold all possible entries. This situation includes the following:

- One slot if the option for reserving ID 0 for the boot device is set in IOC Page 8, and
- Two slots if running IR firmware (to hold the two IR volume IDs), and
- Eight slots, which are reserved for the virtual adapter enclosure to hold any direct-attached devices, and
- *N* slots, which are the sum of all reserved mapping entries (one per slot) for all attached enclosures

In the following example, assume you have two enclosures with one slot each and four enclosures with 31 slots each in the topology. You are using IR firmware. Therefore, you need a mapping table with at least 137 entries.

1 + 2 + 8 + (2 * 1) + (4 * 31) = 137 entries

The driver mapping table is created with the number of entries equal to the IOCFacts MaxTargets field. The number of entries is the maximum number of target devices that the firmware can support. If this value, set by NVDATA, is not large enough to hold all of the potential devices for all the attached enclosures, then the driver cannot find available space to reserve IDs for all of the enclosures.

The enclosure configuration information must be modified such that the enclosure reports a number of slots sufficient to hold all potential physical devices plus all SES targets. The slot number assignments fall into the range of StartingSlot through StartingSlot + NumSlots – 1.

5 Expander and Enclosure Changes Required to Implement Enclosure/Slot Mapping Mode

The LSI controllers have two enclosure modes, managed and unmanaged.

5.1 Managed Enclosure Mode

Managed enclosures adhere to strict rules.

1. Populate the Enclosure Logical ID field that resides in the SMP Report General Response Frame. The same ID is assigned to all expanders in the same enclosure.
2. The Starting Connector Element Index must be the same for all expanders in the enclosure.
3. The Max Number of Connector Elements value must be the maximum number of HDD slots that exist in the design plus the number of virtual SCSI enclosure service (SES) ports in the enclosure. For example, for a 60-drive JBOD and three expanders, 63 is the maximum number.
4. All Discover SMP Responses (or Discover List Responses) must have the HDD connector element type set to 0x20 for an HDD receptacle.

   Assign the virtual SES target to 0x2F.
5. The connector element index in LSI IR adapters is the same as the HDD slot number. When you define the SES element layout, place the array device element first in the element list. The LSI IR adapters do not perform SES reading commands, so all information is conveyed through serial management protocols (SMPs) from the expander discover response.
6. Make sure that every time you request the SMP Discover response that you receive a consistent Connector Type and Element Index. Do not allow the response to change dynamically.

Managed enclosures use the MaxTarget resources most efficiently because the enclosure indicates how many devices can be plugged into that enclosure.

5.2 Unmanaged Enclosure Mode

Unmanaged enclosures, like the LSI SAS switch or other midplane switching products that can have a multitude of interconnect options, use the maximum resources of the expander. For example, if the expander has 36 ports, the controller firmware consumes 36 physical targets and however many other virtual ports are defined in the expander. The total is deducted from the maxtarget resource pool even if you do not connect an HDD to a phy.

To act as an unmanaged enclosure, adhere to the following settings.

1. Report General SMP Enclosure Logical ID = 0
2. Connector Element Start Index = 0
3. Max Connector Element Index = 0

6 General Debug Method

Use the following debug methods to debug device mapping issues.

1. From UART, run `pl status` to check if the firmware recognizes the attached targets. If you see the devices, the firmware is functioning.
2. If the firmware shows the targets, but you do not see the targets in the OS (`fdisk -l` in the Linux OS or Disk Management in the Windows OS), the device mapping might be incorrect.

3. As described in Check Your Mapping Mode and Persistence, discover the mapping mode, and act accordingly.
4. Get a SAS trace, check the SMP Report General, check the Discover SMP response, and check if the responses meet the requirements in Section 5, Expander and Enclosure Changes Required to Implement Enclosure/Slot Mapping Mode.


LSI