

To: T10 Technical Committee
 From: Ed D'Avignon, Vitesse Semiconductor Corp. (davignon@vitesse.com),
 Date: 13 June 2005
 Subject: T10/05-236r0 SAS 2.0 Zoning Wide Port Group Assignments

Revision History

Revision0 (7 June 2005) first revision

Related Documents

sas1r09 - Serial Attached SCSI 1.1 revision 9
 T10/05-144r2 SAS-2 Zoning

Overview

Many SAS system topologies rely on wide ports and shared links to allow the host(s) to communicate with the disk drives. In some systems, both high priority and low priority targets must share the same links. With zoning it can be possible to reserve one or more links of a wide port for high priority targets.

Changes required to the SAS 2.0 Zoning proposal (T10/05-144) to allow the reservation of shared links of a wide port are described.

There are three changes to the current SAS 2.0 zoning proposal are necessary:

1. Allow phys in a wide port to be assigned to different groups (this is not permitted nor prohibited in the currently proposal).
2. Permit trusted phys to be in any group, not just group 127
3. Modify the Broadcast Address Frame rules to work with multiple groups in a wide port.

Suggested Changes

4.1 Zoning model

4.1.2 Zoning Configuration

The Expander zoning configuration consists of two parts: the per expander phy configuration, and per expander zoning permission table.

Every expander phy is associated with the following per-phy configuration parameters:

Table 1. Per phy zoning configuration (PHY_ZONE CONFIGURATION)

Name	Description
TRUSTED	<p>If set to 0, this phy is on the boundary of the zoning fabric. All message (primitives and frames) that come across this phy shall be mapped to be backwards-compatible to SAS standard without zoning features, except for the new SMP commands defined by the zoning extension.</p> <p>If set to 1, this phy is inside the fabric boundary. The new primitives and frame formats that are defined by the zoning extension are allowed to pass through this phy.</p>
GROUP ID[6:0]	<p>The GID defines the zoning Group ID in the range from 0..127.</p> <p>GID=0: Group 0 is a special group that is not allowed to communicate with any other group except for group 127. Note that a device belonging to group 0 can still discover all the expanders and communicate with the SMP virtual target in the expanders (i.e. SMP virtual target within the zoning expanders are considered to have GID=127).</p> <p>GID=127: Group 127 is a special group that is allowed to communicate with all other groups. All trusted phys shall be automatically assigned to have GID=127 by the zoning expanders.</p> <p>GID=1..126: User defined groups. The communications amongst the user defined groups are restricted by the zoning permission table.</p>

Name	Description
SUPERVISOR	If Set to 1, the device attached to this phy is allowed to originate SMP commands to set up and change zoning configuration. If set to 0, the device attached to this phy is not allowed to originate SMP commands to change the zoning information.
SOURCE CHECK	This specifies whether the specified phy shall check the SOURCE SAS address against the SAS address in the IDENTIFY address frame received on the specific phy.

Each phy in a wide port may be assigned a different Group ID. At least one phy in each port with trusted phys should be assigned Group ID 127.

All phys in a wide port shall be trusted or untrusted, the mixing of trusted and untrusted phys in a wide port is not permitted.

...

4.1.6 Broadcasts

...

When a zoned expander detects a phy event or received a broadcast primitive that would normally spawn a BROADCAST (CHANGE), it instead sends a ZONED BROADCAST address frame with the SGID field set to the Group ID (GID) of the phy which caused the event, to all phys that are accessible from the source Group.

If the BROADCAST (CHANGE) primitive is received on a phy of a wide port that has multiple Group IDs on its phys, the SGID field in the BROADCAST address frame shall be set to 127 if any phy in the wide port is in Group ID 127. If no phy in the wide port is in Group ID 127, the SGID field in the BROADCAST address frame shall be set to the Group ID of the receiving phy.

...

4.1.8 Zone checks

...

If neither table routing nor direct routing finds a match for the destination and the expander supports subtractive routing, the OPEN is routed according to subtractive routing rules ~~without zoning permission check. If the subtractive port is used to connect to an adjacent expander, it is a requirement that the subtractive port must be trusted and set to group 127. This allows the OPEN to be forwarded to the next expander which has the topology knowledge to do table or direct routing and the permission check. If the zoning permission table does not have a permission between the SGID and DGID of the subtractive phy, the OPEN is rejected with OPEN_REJECT(ZONING VIOLATION). If the zoning permission table allows routing between the SGID and DGID of the subtractive phy, the expander shall make the connection to the subtractive phy.~~

...