To: T10 Technical Committee
From: Ed D'Avignon, Vitesse Semiconductor (davignon@vitesse.com)
Date: 13 June 2005
Subject: T10/05-235r1 SAS 2.0 Allow Zoning Enhancements to Leave Expanders

Revision History
Revision0 (7 June 2005) first revision

Revision1 (15 June 2005)
Add the ZONE AWARE bit to the REPORT GENERAL command
Require devices that are ZONE AWARE to support the CONFIGURE ZONE PERMISSION and REPORT ZONE PERMISSION SMP Commands in their SMP Targets.
Have the Supervising Expander send the CONFIGURE ZONE PERMISSION command to all devices that are ZONE AWARE, not just expanders

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

Overview
T10/05-144r2 SAS 2 Zoning proposes new OPEN Address Frame fields and a new Broadcast Address Frame. Before leaving a PHY which is not marked TRUSTED, the new OPEN Address Frame fields are replaced with Zeros and the Broadcast Address Frame is replaced with a Broadcast Primitive.

It may be desirable in the future for end devices to receive the zone information contained in the OPEN Address Frame or Broadcast Address Frame. Since end devices are generally not trusted, this information would not be available.

A ZONE Aware bit for the PHY and behavior when the ZONE Aware bit is active are proposed.

Note: change bars in this document refer to the text in T10/05-144r2

Suggested Changes

4.1 Zoning model

4.1.2 Zoning Configuration

Table 1. Per phy zoning configuration (PHY_ZONE CONFIGURATION)

Name	Description
ZONE AWARE	If set to 0, this phy is connected to a device which is not aware of the zoning features. All messages (primitives and frames) that are transmitted from this phy shall be mapped to be backwards-compatible to the SAS standard without zoning feature, except for the new SMP commands defined by the zoning extension. If set to 1, this phy is aware of the zoning features. The new primitives and frame formats that are defined by the zoning extension are transmitted from this phy unmodified. New primitives and frame formats received by this phy are possibly modified depending upon the value of the TRUSTED bit.

Name	Description
TRUSTED	If set to 0, this phy is on the ~~boundary of the zoning fabric~~not trusted. All message (primitives and frames) that ~~come across this phy shall be mapped to be backwards-compatible to SAS standard without zoning features~~are received by this phy are modified to contain the Access Zone Management bit and Source Group ID assigned to this phy~~, except for the new SMP commands defined by the zoning extension~~. If set to 1, this phy is ~~inside the fabric boundary~~trusted. The ~~new primitives and frame formats that are defined by the zoning extension are allowed to pass through this phy~~frames and primitives are allowed to pass unchanged.
GROUP ID[6:0]	The GID defines the zoning Group ID in the range from 0..127. 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). 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. GID=1..126: User defined groups. The communications amongst the user defined groups are restricted by the zoning permission table.
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.

4.1.4 OPEN address frame handling

The OPEN address frame used in a zoned SAS environment includes a new ACCESS ZONE MANAGEMENT bit and the SOURCE GROUP ID field (in the COMPATIBLE FEATURES area). The ACCESS ZONE MANAGEMENT bit and SOURCE GROUP ID field in these OPEN address frames are only valid among devices inside the zoning fabric on trusted expander phys.

When an untrusted expander phy receives an OPEN frame, it sets the ACCESS ZONE MANAGEMENT bit according to the value of SUPERVISOR bit of the expander phy, and sets the SOURCE GROUP ID according to the value of the GROUP ID of the expander phy. When a~~n untrusted~~ non Zone Aware expander phy transmits or forwards an OPEN, it sets ACCESS ZONE MANAGEMENT bit and the SOURCE GROUP ID to zero.

When a ~~trusted~~ Zone Aware expander phy transmits or forwards an OPEN, the value of the ACCESS ZONE MANAGEMENT bit and the SOURCE GROUP ID are preserved and transmitted. This mechanism allows the use of the new OPEN frame format inside zoning fabric across ~~trusted~~ Zone Aware phys, and ensures legacy OPEN address frame format is used outside the zoned fabric boundary This preserves backwards compatibility.

…

4.1.5 SMP functions

…

The SMP DISCOVER command is extended to report the ZONE VIOLATION, ~~SOURCE CHECK~~ZONE AWARE, TRUSTED SUPERVISOR and GROUP ID information that is part of the PHY ZONE configuration of the specific Phy. When the SMP DISCOVER command is executed from a source group

(as indicated by the SGID in the OPEN frame that set up the SMP connection), the zoning expander shall report the accurate information for the Phys that the source group is allowed to access according to the ZONE PERMISSION table. The Phys that are inaccessible from the source group shall be reported as VACANT.

…

The SMP REPORT ZONE ROUTE TABLE reports the zone route table, which is an extension of the routing table defined by SAS 1.1. The ZONE route table is logically organized in a similar way as the routing table. Each entry of the table is extended to contain fields in addition to attached SAS address including ZONE AWARE, TRUSTED, SUPERVISOR, GROUP ID, ATTACHED DEVICE TYPE.

4.1.6 Broadcasts

…

When transmitting a ZONED BROADCAST event on a ~~trusted~~ zone aware expander phy, the expander transmits a ZONED BROADCAST address frame and sets the outgoing SGID field to the SGID field received (or the phy GID of the phy causing the broadcast).

A BROADCAST primitive received on a trusted phy shall be treated the same way as a BROADCAST address frame with SGID of 127 (unrestricted broadcast).

A BROADCAST primitive received on a untrusted phy shall be treated the same way as a BROADCAST address frame with the SGID assigned to the Group ID of the receiving phy.

A BROADCAST address frame received on an untrusted phy shall have its SGID changed to the Group ID of the receiving phy.

When transmitting a ZONED BROADCAST event on a ~~untrusted~~ non zone aware phy, the expander transmits a BROADCAST primitive that represents the type of broadcast event represented by the ZONED BROADCAST event, but the SGID information is discarded.

…

10.4.3 SMP functions

10.4.3.3 REPORT GENERAL function

…

Table 167 — REPORT GENERAL response

Byte\Bit	7	6	5	4	3	2	1	0
0	SMP FRAME TYPE (41h)							
1	FUNCTION (10h)							
2	FUNCTION RESULT							
3	Reserved							
4	(MSB)	EXPANDER CHANGE COUNT						
5								(LSB)
6	(MSB)	EXPANDER ROUTE INDEXES						
7								(LSB)
8	Reserved	NUMBER OF ZONES						
9	NUMBER OF PHYS							
10	Reserved			ZONE AWARE	Reserved		CONFIGURING	CONFIGURABLE ROUTE TABLE
11	Reserved							
12	ENCLOSURE LOGICAL IDENTIFIER							
19								
20	Reserved							
27								
28	(MSB)	CRC						
31								(LSB)

The NUMBER OF ZONES field indicates the number of zones supported ~~when the ZONING SUPPORTED bit is 1~~. Expander devices that support zoning shall support this field. Other device types may not support this field. For expanders that do not support zoning, this field should be set to 0. Note that group 0 and group 127 must be supported in all zoning expanders. The remaining zone indexes should range from 1 to (NUMBER OF ZONES –2).

The ZONE AWARE bit indicates whether the device supports the zoning feature and understands zoning constructs such as BROADCAST frames, the SGID of the OPEN ADDRESS frame and the zoning SMP commands.

…

10.4.3.5 DISCOVER function

….

Table 171 defines the response format.

Table 171 — DISCOVER response

Byte\Bit	7	6	5	4	3	2	1	0
0	SMP FRAME TYPE (41h)							
1	FUNCTION (10h)							
2	FUNCTION RESULT							
3	Reserved							
4	Ignored							
7								
8	Reserved							
9	PHY IDENTIFIER							
10	Ignored							
11	Reserved							
12	Ignored	ATTACHED DEVICE TYPE			Ignored			
13	Reserved				NEGOTIATED PHYSICAL LINK RATE			
14					ATTACHED SSP INITIATOR	ATTACHED STP INITIATOR	ATTACHED SMP INITIATOR	ATTACHED SATA HOST
15	ATTACHED SATA PORT SELECTOR	Reserved			ATTACHED SSP TARGET	ATTACHED STP TARGET	ATTACHED SMP TARGET	ATTACHED SATA DEVICE
16	SAS ADDRESS							
23								
24	ATTACHED SAS ADDRESS							
31								
32		ATTACHED PHY IDENTIFIER						
33	Reserved							
39								
40	PROGRAMMED MINIMUM PHYSICAL LINK RATE				HARDWARE MINIMUM PHYSICAL LINK RATE			
41	PROGRAMMAED MAXIMUM PHYSICAL LINK RATE				HARDWARE MAXIMUM PHYSICAL LINK RATE			
42	PHY CHANGE COUNT							
43	VIRTUAL PHY	Reserved			PARTIAL PATHWAY TIMEOUT VALUE			
44	TRUSTED	ZONE VIOLATION	SOURCE CHECK	Reserved	ROUTING ATTRIBUTE			
45	Reserved	CONNECTOR TYPE						
46	CONNECTOR ELEMENT INDEX							
47	CONNECTOR PHYSICAL LINK							
48	Reserved				ZONE VIOLATION	~~SOURCE CHECK~~ ZONE AWARE	TRUSTED	SUPERVISOR
49	Reserved	GROUP ID						
50	Vendor Specific							
51								
52	(MSB)	CRC						
55								(LSB)

…

The ZONE VIOLATION field is set to 1 if any ZONE violation has occurred causing the specified phy to send OPEN_REJECT(ZONE VIOLIATION). The ZONE VIOLATION shall be cleared if a PHY CONTROL function with operation code of CLEAR ERROR LOG for the specified phy is received from a supervisor.

The TRUSTED bit reports whether the specified phy is currently configured as trusted phy or untrusted phy by the supervisor.

The SUPERVISOR bit reports whether the specified phy is currently configured as a zone supervisor phy.

The ZONE AWARE bit reports whether the device attached to the specified phy understands the zoning extension to the OPEN Address Frame , the BROADCAST Address Frame and the zoning SMP commands.

The GROUP ID fields reports the source group ID assignment of the specified phy.

…

10.4.3.12 CONFIGURE PHY ZONE function

…

Table 7 defines the PHY ZONE configuration entry descriptor.

Table 7 — PHY ZONE configuration entry descriptor

Byte\Bit	7	6	5	4	3	2	1	0
0						~~SOURCE CHECK~~ Zone Aware	TRUSTED	SUPERVISOR
1	Reserved	GROUP ID						

The GROUP ID field specifies the group ID to be assigned to the specified phy.

The SUPERVISOR field specifies whether the specified phy is a supervisor.

The TRUSTED field specifies whether the specified phy is trusted or untrusted.

The ZONE AWARE bit reports whether the device attached to the specified phy understands the zoning extension to the OPEN Address Frame and BROADCAST Address Frame.

….

10.4.3.13 CONFIGURE ZONE PERMISSION function

…

This SMP function shall be supported by SMP target ports in ~~expander~~ devices if the ~~ZONING SUPPORTED~~ZONE AWARE bit is set to one in the REPORT GENERAL function. Devices with the ZONE AWARE bit set to zero shall not support this command. The SMP target should only execute the CONFIGURE ZONE PERMISSION function if the OPEN frame that set up the SMP connect has the ACCESS ZONE MANAGEMENT bit set to 1. If the SMP OPEN frame has ACCESS ZONE MANAGEMENT bit set to 0, the SMP target should generate a SMP response frame with FUNCTION RESULT set to UNKNOWN SMP FUNCTION.

…

If a non-supervising ~~expander~~ device receive~~s~~ a CONFIGURE ZONE PERMISSION command with PROPAGATE UPDATE =0, the non-supervising ~~expander~~ device shall update~~s~~ its own permission table without propagation.

If a non-supervising ~~expander~~ device receives a CONFIGURE ZONE PERMISISON command with PROPAGATE UPDATE =1, the non-supervising ~~expander~~ device shall ignore this command and return FUNCTION FAILED.

If a supervising expander receives CONFIGURE ZONE PERMISSION command with PROPAGATE UPDATE =1, the supervising expander shall update its permission table accordingly and propagate the same permission table updates to all other ~~zoning expanders~~zone aware devices within the domain by sending SMP CONFIGURE ZONE PERMISSION commands to individual ~~expanders~~ devices with PROPGATE UPDATE =0. If the supervising expander receives another CONFIGURE ZONE PERMISSION command before the propagate of the zone permission update from the previous command has not been completed, the supervising expander shall ignore the new CONFIGURE ZONE PERMISSION command and return FUNCTION FAILED.

…

10.4.3.14 REPORT ZONE PERMISSION function

The REPORT ZONE PERMISSION function reports the zoning permission table of the ~~expander~~device.

This SMP function shall be supported by SMP target ports in ~~expander~~ devices if the ~~ZONING SUPPORTED~~ ZONE AWARE bit is set to one in the REPORT GENERAL function. The SMP target should only execute the CONFIGURE ZONE PERMISSION function if the OPEN frame that set up the SMP connect has the ACCESS ZONE MANAGEMENT bit set to 1. If the SMP OPEN frame has ACCESS ZONE MANAGEMENT bit set to 0, the SMP target should generate a SMP response frame with FUNCTION RESULT set to UNKNOWN SMP FUNCTION.

…

10.4.3.15 REPORT ZONE ROUTE TABLE function

…

Table 20 defines the ZONE route table entry descriptor.

Table 20 —ZONE ROUTE TABLE entry descriptor

Byte\Bit	7	6	5	4	3	2	1	0
0	DISABLE EXPANDER ROUTE ENTRY	Reserved						
1	Ignored	ATTACHED DEVICE TYPE			~~Ignored~~	ZONE AWARE	TRUSTED	SUPERVISOR
2	Ignored	GROUP ID						
3	Ignored							
4	ROUTED SAS ADDRESS							
11								

The DISABLE EXPANDER ROUTE ENTRY bit specifies whether the ECM shall use the expander route entry to route connection requests (see 4.6.7.3). If the DISABLE EXPANDER ROUTE ENTRY bit is set to zero, then the ECM shall use the expander route entry to route connection requests. If the DISABLE EXPANDER ROUTE ENTRY bit is set to one, the ECM shall not use the expander route entry to route connection requests.

The SUPERVISOR field specifies whether the specified SAS address corresponds to a supervisor.

The TRUSTED field specifies whether the specified SAS address is trusted or untrusted.

The ZONE AWARE bit reports whether the device attached to the specified phy understands the zoning extension to the OPEN Address Frame and BROADCAST Address Frame.

The ROUTED SAS ADDRESS field contains the routed SAS address for the expander route entry being configured (see 4.6.7.3).

The GROUP ID field contains the GROUP ID for the expander route entry being configured (see 4.6.7.3).

The ATTACHED DEVICE TYPE field indicates the DEVICE TYPE value received during the link reset sequence and is defined in table 178.