

To: T10 Technical Committee
From: Tim Hogle, LSI Logic (tim.hogle@lsil.com)
Date: January 24, 2006
Subject: T10/06-048r1 SAS-2 zoning enhancements

Revision Information

- Revision 0 Initial
- Revision 1 Incorporated feedback based on January 9, 2006 SAS Protocol Working Group discussions
 - remove management and broadcast references
 - updated usage model diagram for bridge device (Figure 5)
 - added example diagram (Figure 6)
 - update Table 2 and clarify how source group is determined as a function of the routing attribute of the receiving phy

References

SAS2r01 Serial Attached SCSI - 2 (SAS-2)
T10/06-019r1 SAS-2 zoning
T10/06-029r2 SAS-2 Expander Configuration Supervisor and SMP CHANGE Request

Introduction

Reference document T10/06-019 provides a baseline view of zoning functionality for a SAS service delivery subsystem. However, the proposed view of zoning is not sufficiently broad to encompass expected usage models for the SAS-2 market timeframe.

This proposal seeks to define zoning methods for SAS-2 which satisfy expected usage models by building on certain aspects of the method proposed by T10/06-019.

~~T10/06-019 defines centralized management of the zoning service delivery subsystem, however zone management is defined independently from expander self-configuration. This proposal seeks to align zone management with expander self-configuration.~~

T10/06-019 defines zone groups and a zone permission table to establish access permission (whether a connection is allowed between source zone group and a destination zone group). Zone groups and the zone permission table provide a good method of providing zoning functionality and this proposal use them as the basis for SAS-2 zoning methods.

In order to satisfy a broader scope of usage models, the primary difference between this proposal and T10/06-019 is how source group assignment is determined. Whereas T10/06-019 assigns zone groups based **solely** on a phy attribute, this proposal allows the **source** zone groups to be **determined via a zone route table lookup. assigned by SAS address.** Within this proposal, **phy-resolved source zone group mapping shall mean the method provided by T10/06-019 and address-resolved source zone group mapping shall mean the enhanced method offered by this proposal.**

Usage model discussion – establish required scope of SAS-2 zoning

Usage model 1: Traffic segregation as illustrated by Figure 1.

Topology/resource partitioning in the manner of VLAN/VSAN

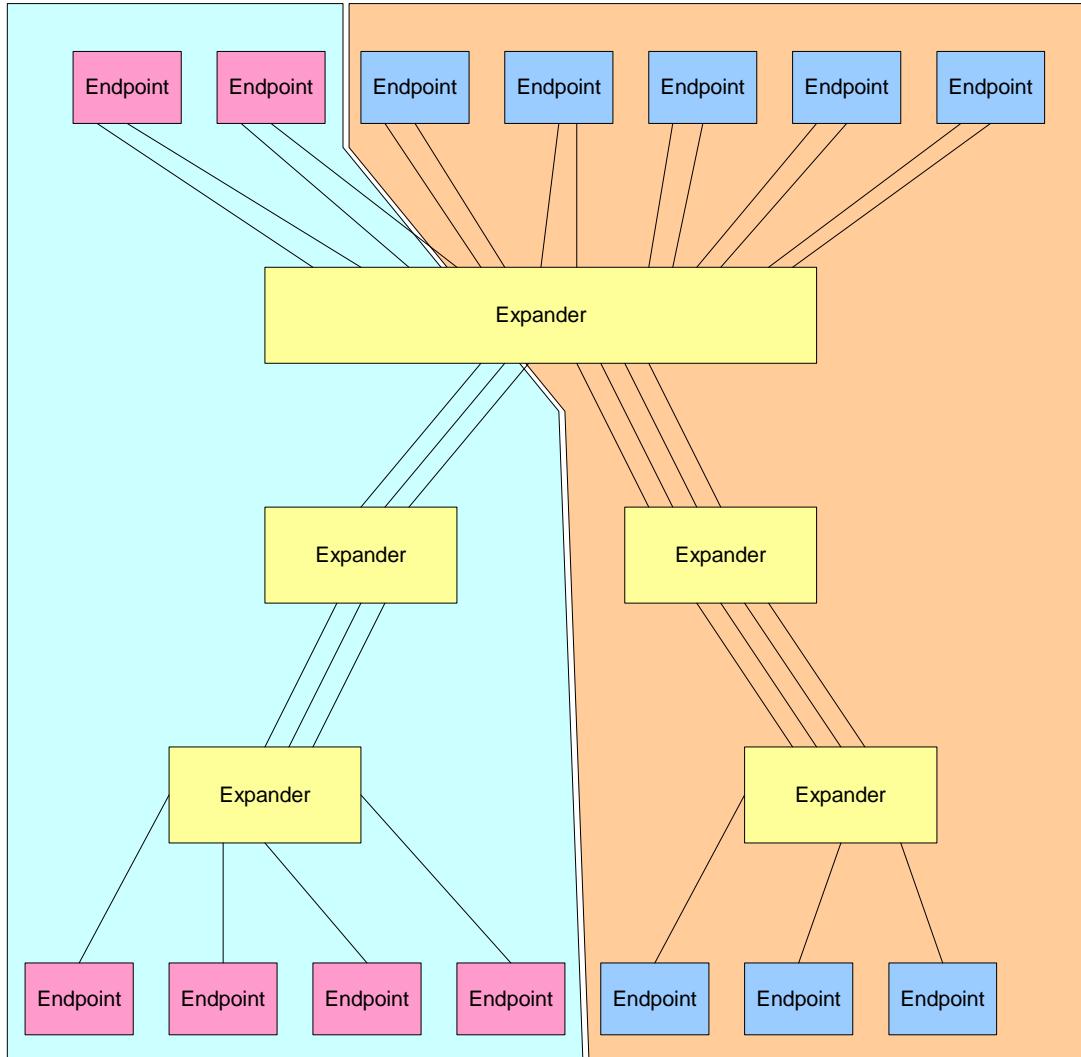


Figure 1: Traffic Segregation

Either phy-resolved or address-resolved source zone group mapping can satisfy the traffic segregation usage model.

However, phy-resolved source zone group mapping cannot differentiate traffic when legacy (non-zoned) expanders are present within the topology.

Usage model 2: Access control as illustrated by Figure 2.

Allow control of which hosts can access which devices.

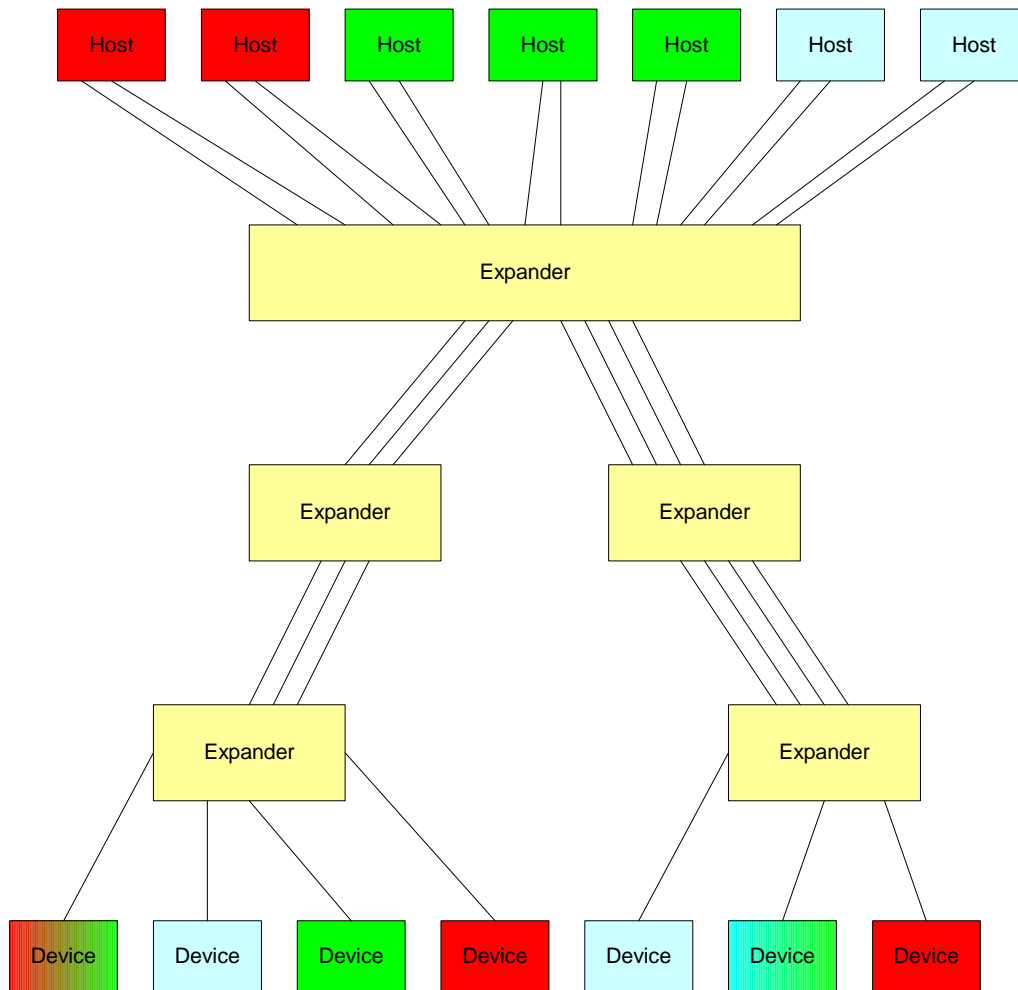


Figure 2: Access control

Either phy-resolved or address-resolved source zone group mapping can satisfy the Access control usage model.

However, phy-resolved source zone group mapping cannot provide full access control when legacy (non-zoned) expanders are present within the topology.

Usage model 3: Device sharing as illustrated by Figure 3.
Hosts sharing access to devices to not allowed to see or interfere with each other.

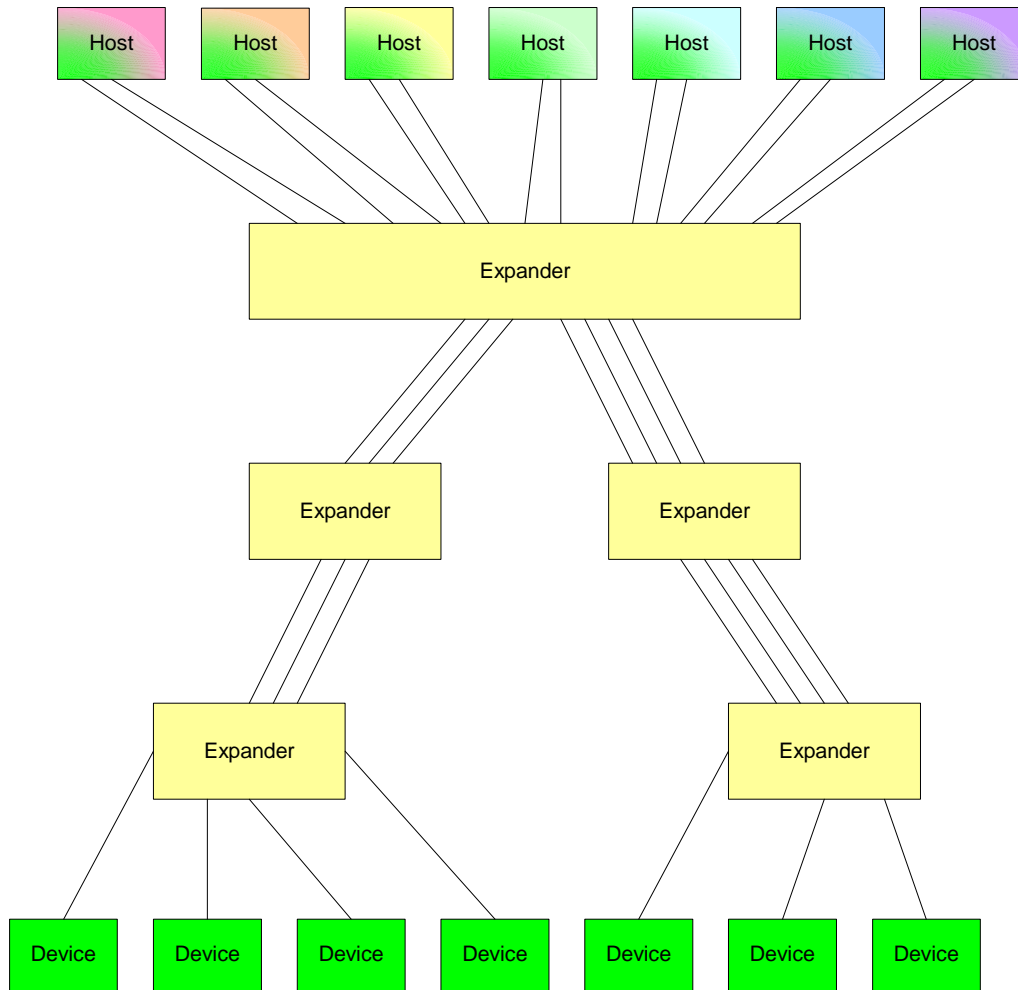


Figure 3: Device sharing

Either phy-resolved or address-resolved source zone group mapping can satisfy the Device sharing usage model.

Usage model 4: Legacy (non-zoned) JBODs as illustrated by Figure 4.
Adding zoned root expanders that resolve source group information from a SAS address allows complete zoning functionality to be realized without changes to the JBODs.

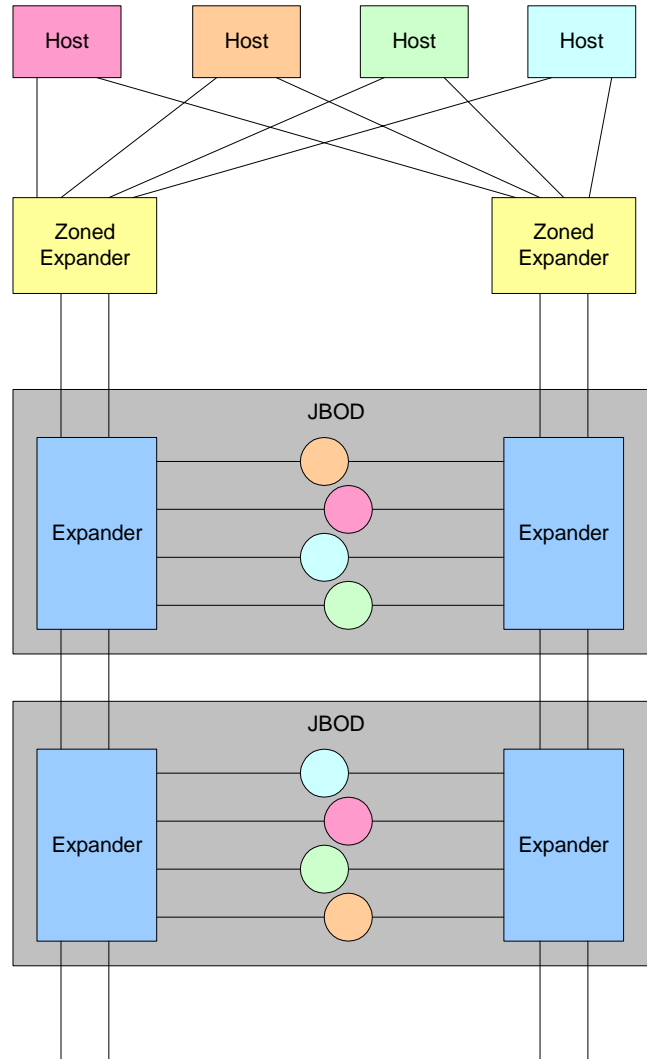


Figure 4: Legacy JBODs

This usage model cannot be realized with the phy-resolved source zone group mapping.

Usage model 5 Bridge or Virtual Devices as illustrated by **Figure 5**. Providing the ability to properly differentiate zoning information on behalf of bridged devices, e.g. **Fibre Channel to SAS Bridge Device presenting itself as a non-zoning expander with virtual devices**.

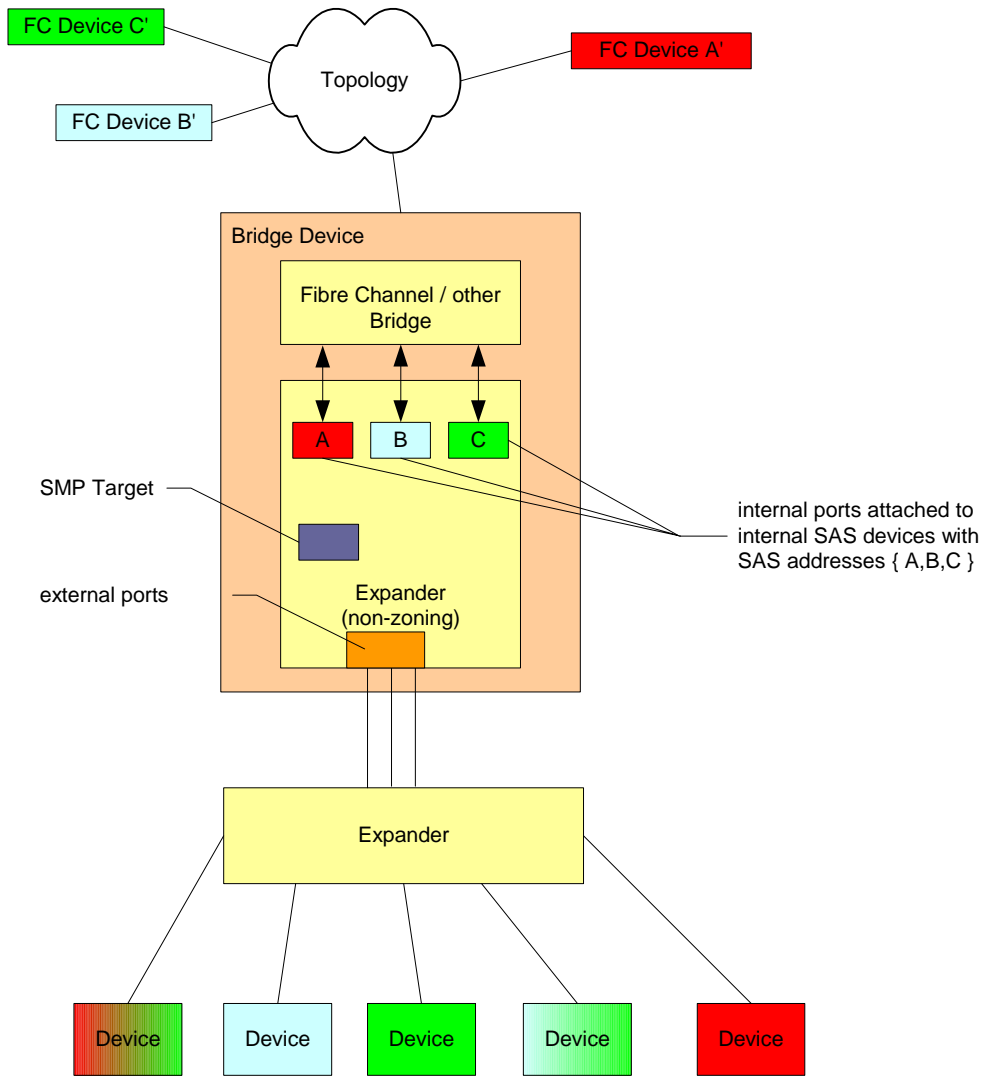


Figure 5: Bridged devices

This usage model cannot be realized with the phy-resolved source zone group mapping.

Address-resolved source zone group mapping example

Figure 6 shows how Address resolved source zone group mapping satisfies the requirements of Usage model 4 Legacy (non-zoned) JBODs.

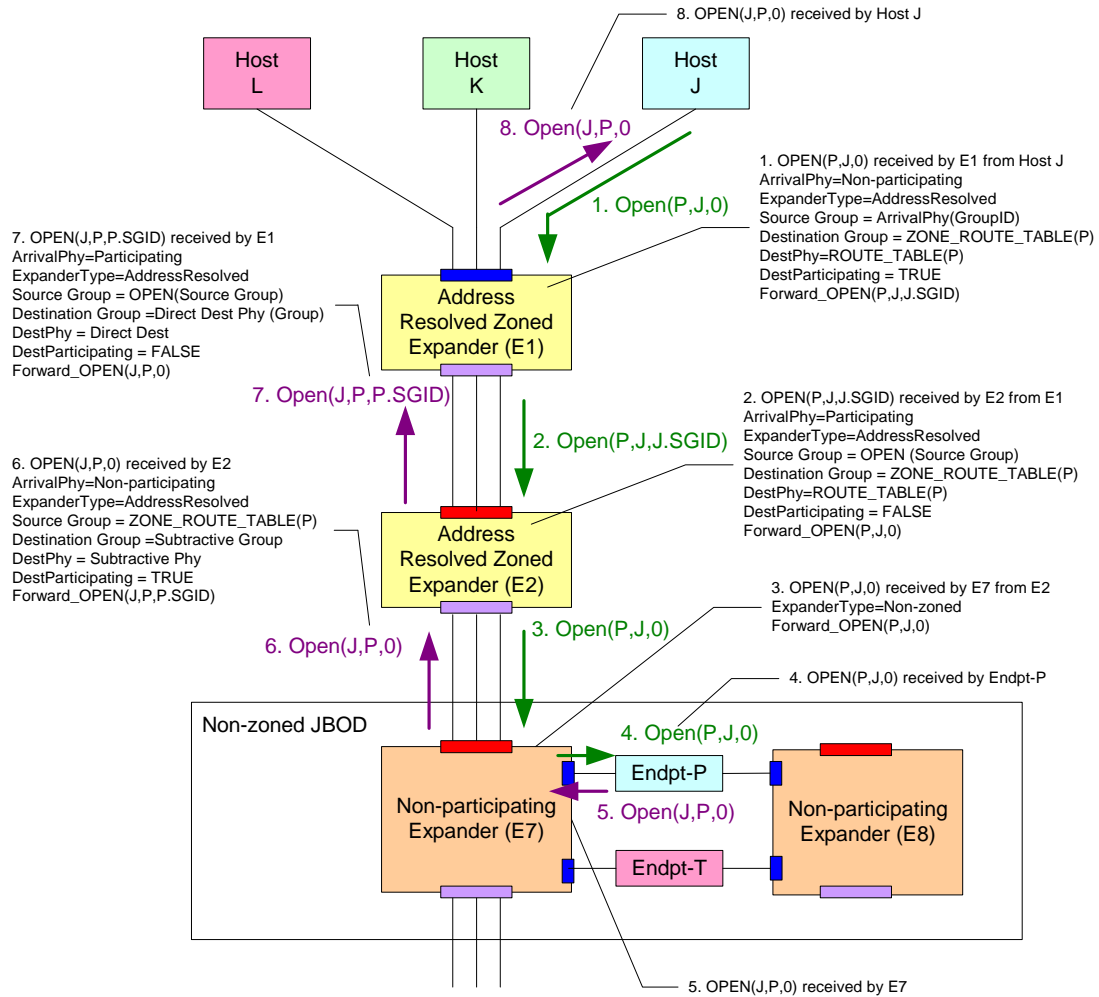


Figure 6: Address-resolved source zone group mapping example

4.9.3.5 Zone Routing

When a zoning expander device receives an OPEN request, the zoning permission table shall check the access permission between the source zone group and the destination zone group. If the zone permission table entry is set to one then access between the phys is allowed and the zoning expander shall continue with the normal ECM arbitration procedure. If the zone permission table entry is set to zero then access is not permitted and the OPEN_REJECT (ZONE VIOLATION) response shall be sent for the OPEN request and the ZONE VIOLATION bit for the source phy shall be set to one.

The zoning expander device uses the rules in Table 1 and Table 2 to check the zone group access permission of the OPEN request.

Table 1 defines how the source zone group is determined from a received OPEN address frame.

Table 1 – Source zone group mapping

Within or on boundary of zoned service delivery subsystem	Source zone group mapping			
	Phy-resolved	Address-resolved (route attribute of receiving phy)		
		Direct Attach	Subtractive Routed	Table routed
Within zoned service delivery subsystem (ZONE PARTICIPATING = 1)	Source zone group in OPEN address frame	Source zone group in OPEN address frame	Source zone group in OPEN address frame	Source zone group in OPEN address frame
On boundary of zoned service delivery subsystem (ZONE PARTICIPATING = 0)	Zone group of the receiving phy	Zone group of the receiving phy	Zone group of the receiving phy	Zone group stored in the zone route table for the source SAS address.

Note:

Address-resolved source group mapping differs from Phy-resolved source group mapping only when an OPEN address frame is received on the boundary of the zoned service delivery subsystem and the routing attribute of the receiving phy is table routed.

To provide the broadest scope of potential zoning usage models, Address-resolved zoning expanders should not utilize subtractive phys on the boundary of the zoned service delivery subsystem.

Table 2 defines how the destination zone group is determined from a received OPEN address frame.

Table 2 – Destination zone group mapping

Expander routing attribute	Destination zone group mapping	
	Phy-resolved	Address-resolved
direct routing	Zone group of the destination phy	Zone group of the destination phy
table routing	Zone group stored in the zone route table for the destination SAS address.	Zone group stored in the zone route table for the destination SAS address.
subtractive routing	Zone group of the subtractive phy.	Zone group of the subtractive phy

4.9.5 Access Zone Management

WG discuss – what purpose does the Access Zone Management bit provide? Suggested alternative to Table 10 below (note list of SMP functions is incomplete and should be expanded):

Table 10 - SMP zone function result field responses

Active zone supervisor	OPEN-address-frame ACCESS-ZONE-MANAGEMENT-bit	Zone permission table entry	DISCOVER ^a	REPORT ROUTE INFORMATION	REPORT GENERAL	REPORT MANUFACTURER INFORMATION	REPORT ZONE PERMISSION	REPORT ZONE ROUTE TABLE	CONFIGURE ZONE PERMISSION	CONFIGURE PHY ZONE	
YES	1	1 or 0	SMP FUNCTION ACCEPTED								
NO	0	0 ^a	PHY VACANT	SMP FUNCTION ACCEPTED		UNKNOWN SMP FUNCTION					
NO	0	1	SMP FUNCTION ACCEPTED				UNKNOWN SMP FUNCTION				
NO	1	0 ^a	PHY VACANT	SMP FUNCTION ACCEPTED				SMP FUNCTION FAILED			
NO	1	1	SMP FUNCTION ACCEPTED						SMP FUNCTION FAILED		

a. If the NO ZONE MASK bit is set to one in the DISCOVER request frame then the zone permission table entry is ignored and the response shall report all phy connections.

10 Application Layer...

10.4.3 SMP functions

10.4.3.1 REPORT GENERAL function

Changes to 06-019r1 and 06-029r2 REPORT GENERAL response – Add an ADDRESS RESOLVED ZONE DEVICE bit to indicate that the device is capable of Address-resolved source group mapping.

Table 193 defines the response format.

Table 197 - REPORT GENERAL response

Byte\Bit	7	6	5	4	3	2	1	0
	...							
10	ZONE DEVICE	ADDRESS RESOLVED ZONE DEVICE	Reserved			ACTIVE SUPERVISOR CONFIGURABLE ROUTE TABLE	CONFIGURING	CONFIGURABLE ROUTE TABLE
11	ACTIVE ZONE SUPERVISOR PRIORITY Reserved				ZONE SUPERVISING SUPERVISOR PRIORITY			
12	ENCLOSURE LOGICAL IDENTIFIER							
19	ENCLOSURE LOGICAL IDENTIFIER							
20	ENCLOSURE LOGICAL IDENTIFIER							
27	ACTIVE ZONE SUPERVISOR SAS ADDRESS							
	...							

...

The ADDRESS RESOLVED ZONE DEVICE bit shall be set to one if the device supports Address-resolved source zone group mapping (see 4.9.3.5). The ADDRESS RESOLVED ZONE DEVICE bit shall be ignored if the ZONE DEVICE bit is set to zero.

...

10.4.3.5 DISCOVER function

Changes to 06-019r1 and 06-029r2 DISCOVER response – Add ZONE ADDRESS RESOLVED bit to indicate that a phy within a zoned device is configured for Address-Resolved Source Group Mapping.

Table 197 defines the response format.

Table 197 - DISCOVER response

Byte/Bit	7	6	5	4	3	2	1	0
	...							
33	Reserved			ATTACHED ZONE DEVICE	ATTACHED ZONE BROADCAST METHOD Reserved			
	...							
48	Reserved		ZONE VIOLATION	ZONE PARTICIPATING	ZONE ADDRESS RESOLVED	ZONE SUPERVISING PRIORITY Reserved		
49	Reserved	ZONE GROUP						
	...							

...

The ATTACHED ZONE DEVICE bit indicates the value of the ZONE DEVICE bit received in the IDENTIFY address frame (see 7.8.2) during the identification sequence.

The ATTACHED ZONE BROADCAST METHOD field indicates the value of the ZONE BROADCAST METHOD field (See **Error! Reference source not found.**) received in the IDENTIFY address frame (see 7.8.2) during the identification sequence.

A ZONE VIOLATION bit set to one indicates that the phy sent an OPEN_REJECT (ZONE VIOLATION) response to a prior DISCOVER command. (See section 7.2.5.11)

A ZONE PARTICIPATING bit set to one indicates that the phy is attached to another zoning device.

A ZONE ADDRESS RESOLVED bit set to one indicates that the phy within a zoning expander device is configured to use an Address-resolved method as specified in Table 2 to determine the source group for received OPEN address frames.

A ZONE ADDRESS RESOLVED bit set to zero indicates that the phy within a zoning expander device is configured to use a Phy-resolved method as specified in Table 2 to determine the source group for received OPEN address frames.

The ZONE SUPERVISING PRIORITY field indicates the zone supervisor device election priority of the phy.

The ZONE GROUP field indicates the zone group that contains the phy.

10.4.3.13 CONFIGURE PHY ZONE function

Changes to 06-019r1 and 06-029r2 – Add a ZONE ADDRESS RESOLVED bit to configure a phy within a zoned device for Address-Resolved Source Group Mapping.

Table 18 - Phy zone configuration entry descriptor

Byte\Bit	7	6	5	4	3	2	1	0
0	Reserved			ZONE PARTICIPATING	ZONE ADDRESS RESOLVED	ZONE SUPERVISING PRIORITY Reserved		
1	Reserved	ZONE GROUP						

A ZONE PARTICIPATING bit set to one indicates that the phy is attached to another zoning device. A ZONE PARTICIPATING bit set to zero indicates that the phy is not attached to a zoning device.

A ZONE ADDRESS RESOLVED bit set to one indicates that the phy within a zoning expander device shall use an Address-resolved method as specified in Table 2 to determine the source group for received OPEN address frames.

A ZONE ADDRESS RESOLVED bit set to zero indicates that the phy within a zoning expander device shall use a Phy-resolved method as specified in Table 2 to determine the source group for received OPEN address frames.

~~The ZONE SUPERVISING PRIORITY field is defined in section 10.4.3.5.~~

The ZONE GROUP field is defined in section [10.4.3.5](#).

10.4.3.14 CONFIGURE ZONE PERMISSION function

Use 06-029r2 definition (removes ZONE SUPERVISING PRIORITY and UPDATE PRIORITY).

10.4.3.15 REPORT ZONE PERMISSION function

Use 06-029r2 definition (removes ZONE SUPERVISING PRIORITY).

10.4.3.16 REPORT ZONE ROUTE TABLE function

WG discuss – why are ATTACHED DEVICE TYPE and ZONE PARTICIPATING fields included in the Zone route table entry descriptor?

10.4.3.XX CONFIGURE ZONE ROUTE TABLE function

WG discuss – why is there no CONFIGURE ZONE ROUTE TABLE function?