

To: T10 Technical Committee
From: Tim Hogleund, LSI Logic (tim.hogleund@lsil.com)
Date: January 7, 2006
Subject: T10/06-048r0 SAS-2 zoning enhancements

Revision Information

- Revision 0 Initial

References

SAS2r01 Serial Attached SCSI - 2 (SAS-2)
T10/06-048r0 SAS-2 zoning enhancements T10/06-019r1 SAS-2 zoning
T10/06-029r0 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 on a phy attribute, this proposal allows zone groups to be assigned by SAS address.

Summary Direction for zoning (wrt T10/06-019)

Keep the following:

- Zone groups, zone permission table
- Centralized management but using T10/06-029 methods

Add/Enhance the following:

- Zone group assignment
- IDENTIFY address frame and SMP requirements for zoning management

Remove the following:

- ZONED BROADCAST address frame, reference T10/06-029

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

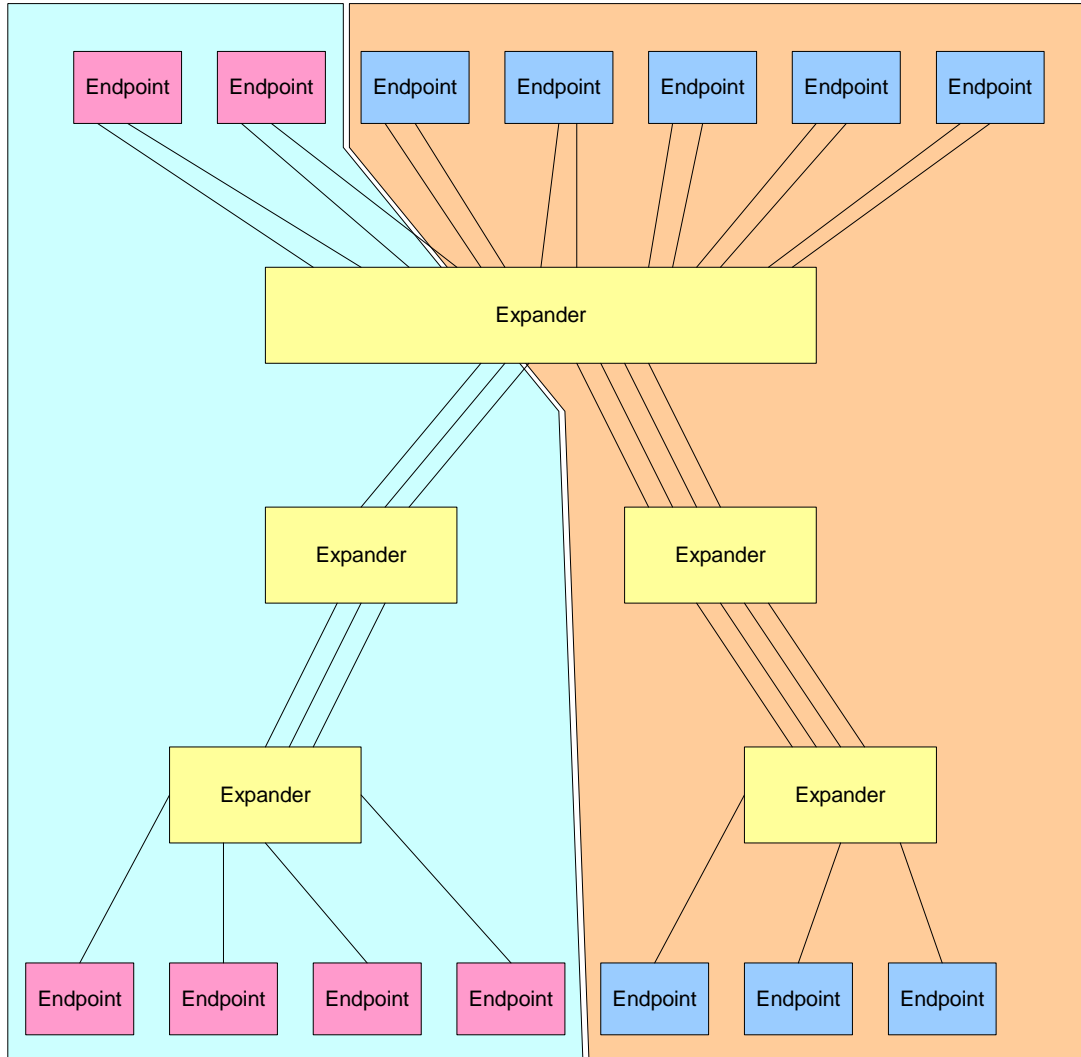


Figure 1: Traffic Segregation

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.
Which hosts can access which devices.

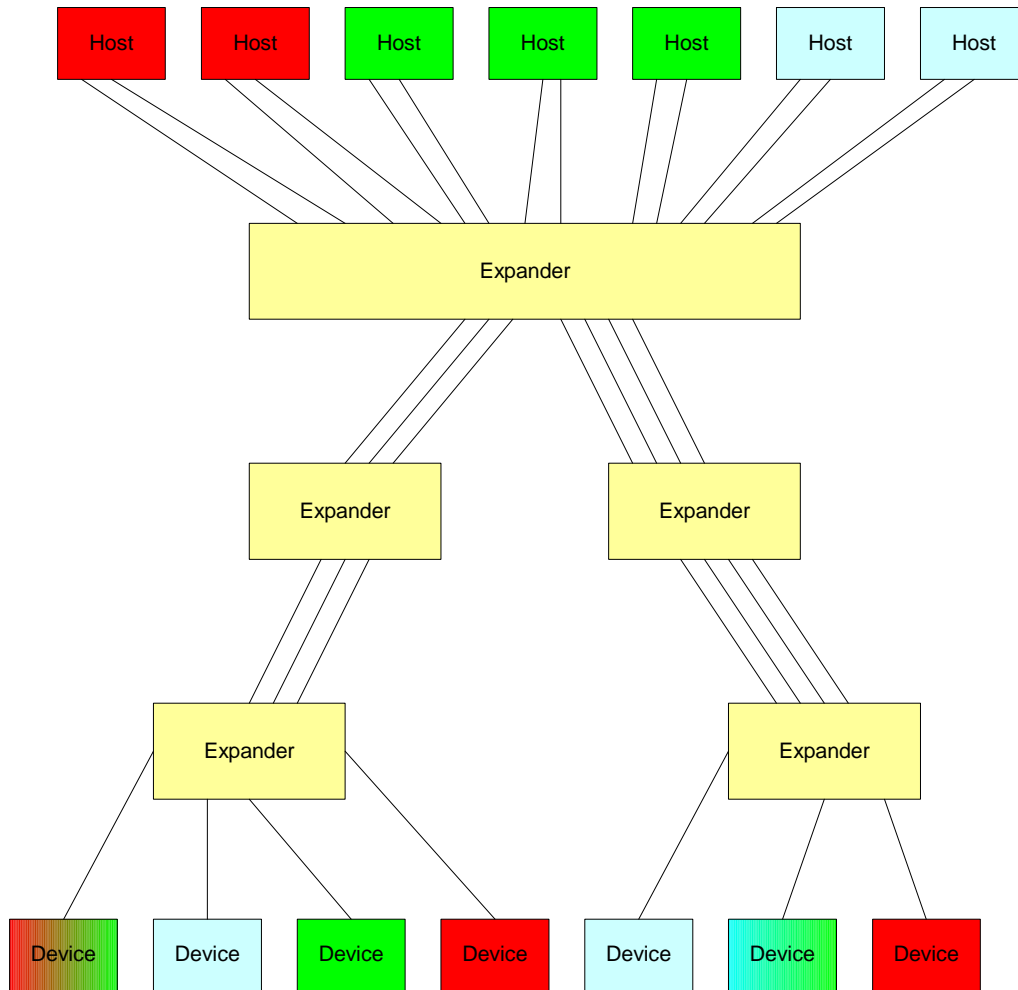


Figure 2: Access control

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.

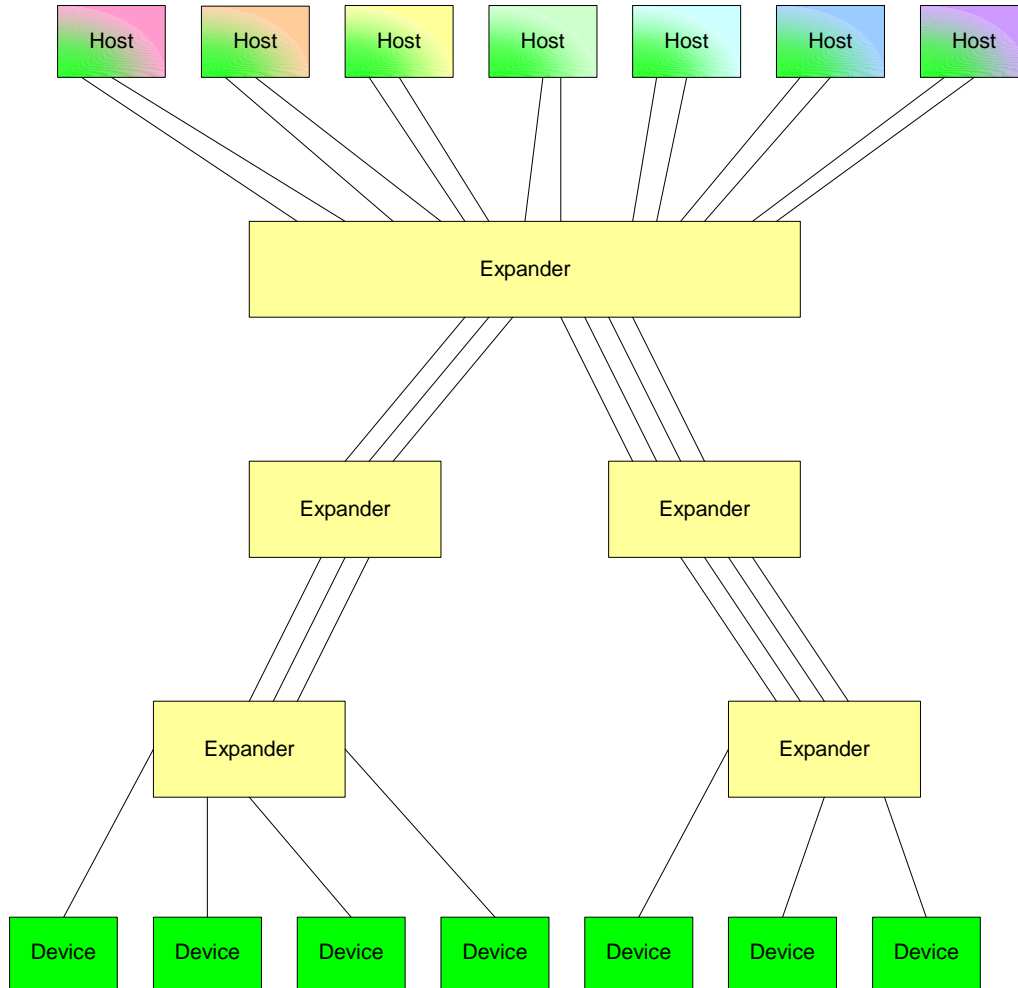


Figure 3: Device sharing

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.

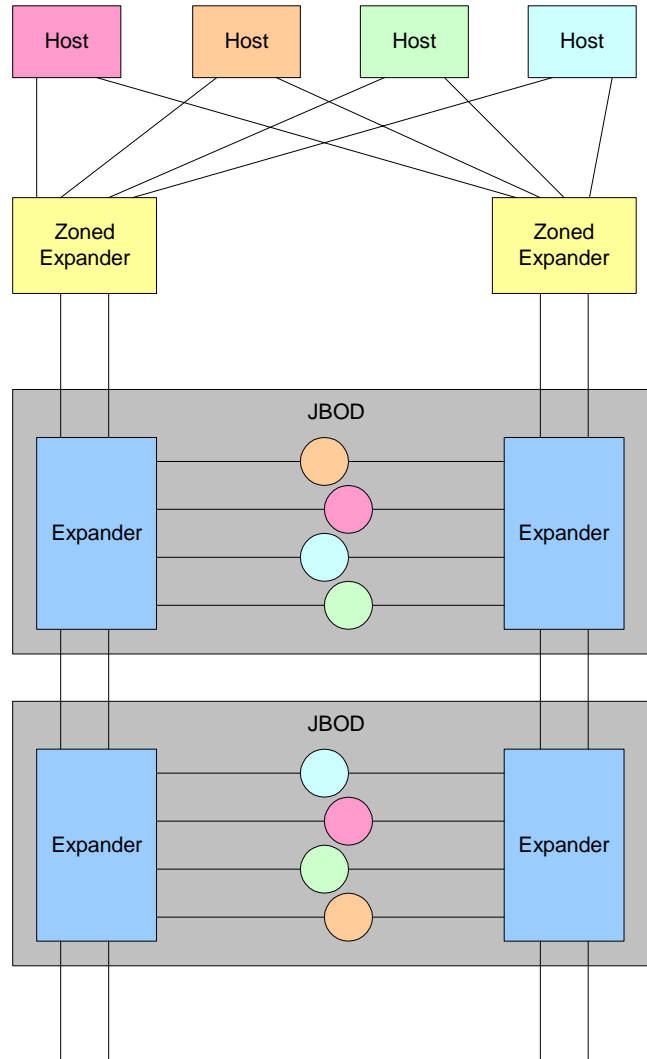


Figure 4: Legacy JBODs

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

Usage model 5 Bridge or Virtual Devices as illustrated by FiN.
Providing the ability to properly differentiate zoning information on behalf of bridged devices.

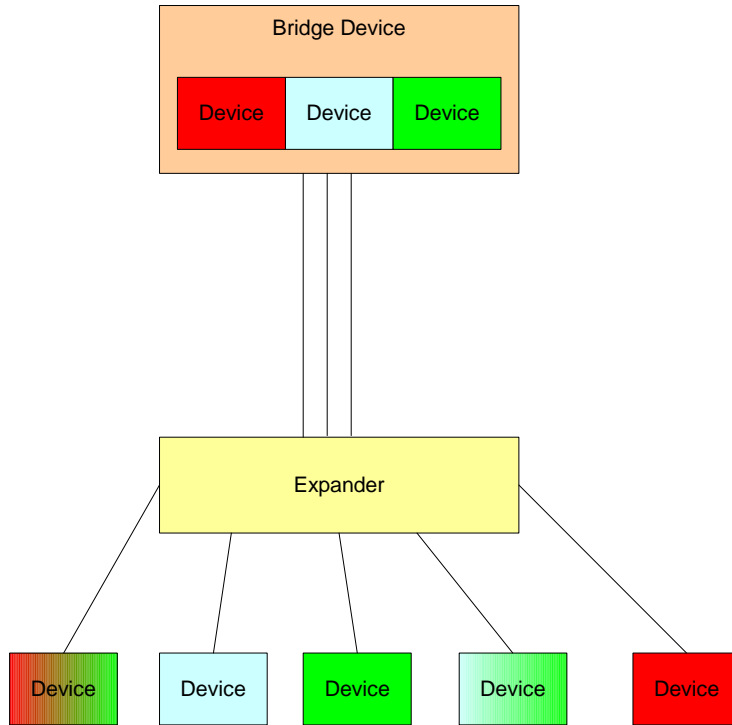


Figure 5: Bridged devices

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

Zone Permission Table

The zone permission table (as defined by T10/06-019) is shown in Table 1.

Table 1 - Zone permission table

| Zone group (X) \ Zone group (Y) | 0 | 1 | 2 | ... | 7 | 8 | ... | 127 |
|---------------------------------|----------------|----------------|-----------------------|-----|-----------------------|-------------------------|-----|--------------|
| 0 | 0 | 1 | 0 | ... | 0 | 0 | ... | 0 |
| 1 | 0 ^a | 1 | 1 | ... | 1 | 1 | ... | 1 |
| 2 | 0 ^a | 1 ^a | Reserved | ... | Reserved | Reserved | ... | Reserved |
| ... | ... | ... | ... | ... | ... | ... | ... | ... |
| 7 | 0 ^a | 1 ^a | Reserved ^a | ... | Reserved | Reserved | ... | Reserved |
| 8 | 0 ^a | 1 ^a | Reserved ^a | ... | Reserved ^a | ZP[8, 8] | ... | ZP[127, 8] |
| ... | ... | ... | ... | ... | ... | ... | ... | ... |
| 127 | 0 ^a | 1 ^a | Reserved ^a | ... | Reserved ^a | ZP[8, 127] ^a | ... | ZP[127, 127] |

^a These fields are ignored for SMP CONFIGURE ZONE PERMISSION request but are reported in SMP REPORT ZONE PERMISSION response. ZP[X, Y] has the same value as ZP[Y, X].

A ZP[X, Y] bit set to one specifies that zone group (X) has permission to access zone group (Y).
A ZP[X, Y] bit set to zero specifies that zone group (Y) has no access zone group (X).

Zone group 0 shall not access any other group except zone group 1. (i.e. ZP[0, 0] and ZP[2.127, 0] shall be set to zero).

Zone group 1 shall access all other zone groups. (i.e. ZP[0.127, 1] shall be set to one).

Zone groups 2 through 7 are reserved zone groups. All reserved bits shall be set to zero. (i.e. ZP[2.7, 2.127] shall be set to zero).

The zone permission table is applicable to both phy-resolved and address-resolved source group methods of zone routing.

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 2 and Table 3 to check the zone group access permission of the OPEN request.

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

Table 2 – Source zone group mapping

| Within or on boundary of zoned service delivery subsystem | Source zone group mapping | |
|---|---|--|
| | Phy based | Address based |
| Within zoned service delivery subsystem | Source zone group in OPEN address frame | Zone group stored in the zone route table for the source SAS address |
| On boundary of zoned service delivery subsystem | Zone group of the receiving phy | Zone group stored in the zone route table for the source SAS address |

Table 3 defines how the destination zone group is determined from a received OPEN address frame.

Table 3 – Destination zone group mapping

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

ZONED BROADCAST address frame

For reasons discussed in T10/06-029, the ZONED BROADCAST address frame shall be removed from SAS-2 zoning and replaced with SMP CHANGE Request.

Zone Management

todo: Define zoning specific SMP management functions

todo: Relate zone management functions to T10/06-029

IDENTIFY address frame

todo: Define IDENTIFY address frame field specifying zoning methods

OPEN address frame

todo: Define how phy-resolved and address-resolved source zone group mapping can be made interoperable. WG discuss.