

T10/06-285r0 SAS-2 Zone configuration Management models.

To: T10 Technical Committee
From: Tim Symons, PMC-Sierra (Tim_Symons@pmc-sierra.com)
Date: 14 June 2006
Subject: 06-285r0 SAS-2 Zone configuration Management models.

Revision Information

- Revision 0 : Models to help evaluate proxy device versus end management device zone configuration.

Referenced Documents

sas2r03a Serial Attached SCSI – 2 (SAS-2) revision 3
06-0-98r5 SAS-2 Self configuring device model
06-215r1 SAS-2 Zone management proxy device (Tim Symons, PMC-Sierra)
06-286r0 SAS-2 Registered Zone management client (Tim Symons, PMC-Sierra)

Overview

Configure zone permission table

Changes in zone permission tables of zoning expander devices is a multi-instruction request to multiple zoning expander devices. A change must be applied to each zoning expander device for consistent behavior and proper operation in a ZPSDS. There is currently no definition within the SAS-2r04 draft for a ZPSDS to ensure that all zoning expander devices have identical zone permission tables. This document provides reference usage models, (based on reference models in document 06-098r5) for discussion of different scenarios representing applications.

Concurrent updates

It is undesirable, but there is a possibility that two independent zone management clients may each issue conflicting zone permission table changes concurrently. Under these circumstances each zoning expander shall have the zone permission table updated, and the “last write” operation shall take effect. In this case both originating zone management clients shall receive “SMP FUNCTION ACCEPTED” response from each zoning expander device, however another zone management client may subsequently overwrite the zone permission table and the change may not result in a notification to all zone management clients (i.e. A zone management client may not be notified of a change)

Under this condition the ZPSDS may have inconsistent zone permission tables.

Zone configuration management

To resolve the risk that a ZPSDS may be partially configured by two zone management clients, with a result of inconsistent zone permission tables in different zoning expander devices, there is a requirement for a definition of zone configuration management.

There are two proposals for zone configuration management;

- a) elected zoning expander proxy device (06-215); or
- b) only one registered zone management client (06-2xx)

Figures 1 and 2 show two zone models for discussion.

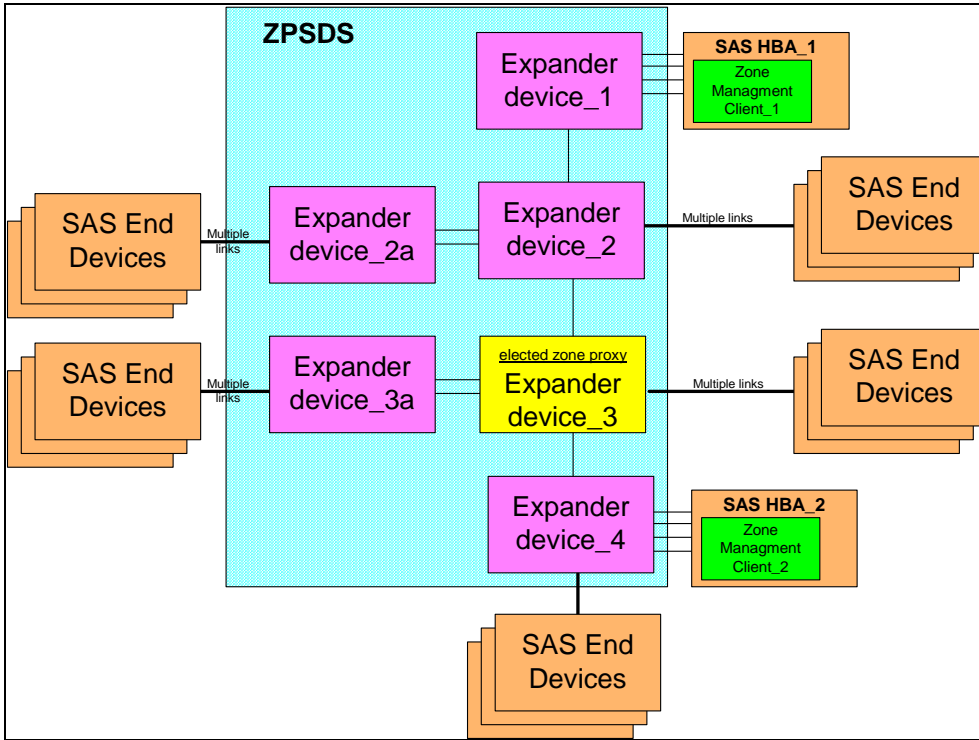


Figure 1 - ZPSDS with elected zoning expander proxy device

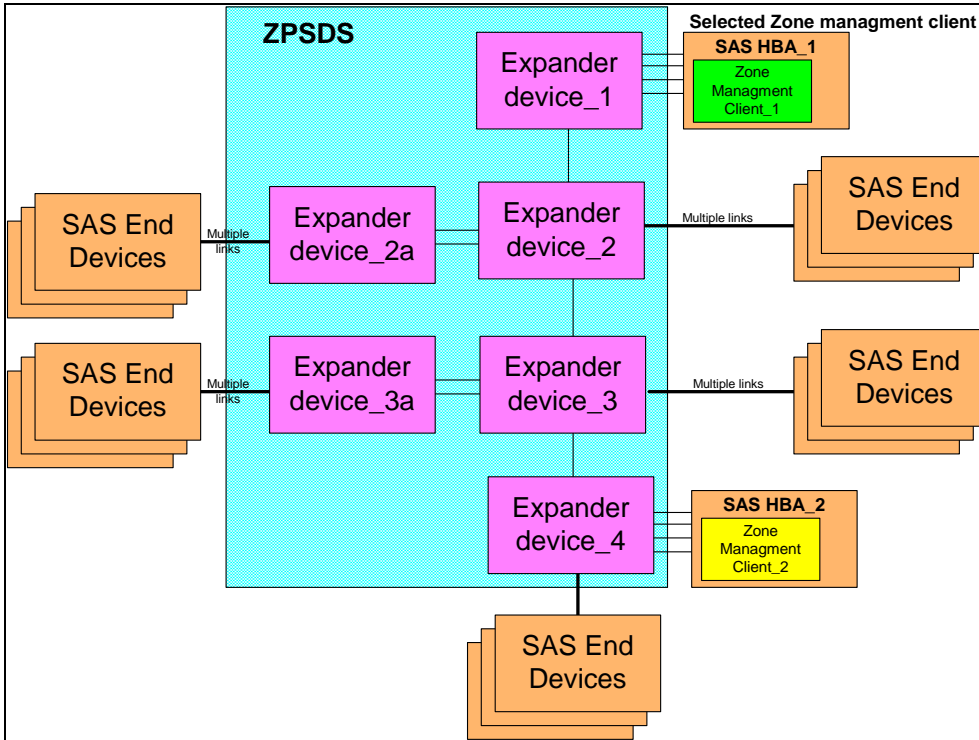


Figure 2 – ZPSDS with a registered zone management client device

Figures 1 and 2 models are assumed for the following walk through examples of configuration events:

Assumptions: Both HBA_1 and HBA_2 have zone access permission for ZG2.

Case 1: HBA 1 issues an SMP CONFIGURE ZONE PERMISSION request

PROXY DEVICE (defined in 06-215):

1. HBA_1 issues SMP CONFIGURE ZONE PERMISSION request to the proxy device, (i.e. zoning expander device_3) identified by the ACTIVE ZONE PROXY SAS ADDRESS field in the REPORT GENERAL response.
2. The proxy device responds with the FUNCTION RESULT field set to ZONE CONFIGURATION ACKNOWLEDGE, indicating that the proxy device successfully received the request and is scheduled to apply the change to all zoning expander devices.
3. The proxy device issues the SMP CONFIGURE ZONE PERMISSION to each zoning expander device in the ZPSDS.
4. Each zoning expander device responds with SMP FUNCTION ACCEPTED or an error condition. Any error condition may result in a retry by the proxy device.
 - a. If all zoning expander device respond with SMP FUNCTION accepted then the zone proxy device sends ZONE CONFIGURATION COMPLETE to HBA_1.
 - b. If any zoning expander device fails to configure then the zone proxy device sends ZONE CONFIGURATION ERROR to HBA_1

REGISTERED ZONE MANAGEMENT CLIENT (defined in 06-286):

Assumptions: HBA_1 is the REGISTERED ZONE MANAGEMENT CLIENT.

1. HBA_1 issues the SMP CONFIGURE ZONE PERMISSION to each zoning expander device in the ZPSDS.
2. Each zoning expander device responds with SMP FUNCTION ACCEPTED or an error condition. Any error condition may result in a retry by the REGISTERED ZONE MANAGEMENT CLIENT

Note: The new configuration is held in a transition buffer, and not effective until the next step.

3. When all zoning expander devices have responded with SMP FUNCTION ACCEPTED, then the registered zone management client shall issue SMP ZONE CONFIGURATION ACTIVATE to all zoning expander devices in the ZPSDS to trigger each zoning expander to apply the updated configuration information. The net result is that the configuration change becomes a monotonic (single event) throughout the ZPSDS, and minimizing the asynchronous nature of a zone permission table update.

Case 2: HBA_2 issues an SMP CONFIGURE ZONE PERMISSION request

PROXY DEVICE (defined in 06-215):

1. HBA_2 issues SMP CONFIGURE ZONE PERMISSION request to the proxy device, (i.e. zoning expander device_3) identified by the ACTIVE ZONE PROXY SAS ADDRESS field in the REPORT GENERAL response.

.... The update proceeds as in case 1.

REGISTERED ZONE MANAGEMENT CLIENT (defined in 06-286):

Assumptions: HBA 1 is the REGISTERED ZONE MANAGEMENT CLIENT.

1. HBA_2 issues the SMP CONFIGURE ZONE PERMISSION to each zoning expander device in the ZPSDS.
2. Each zoning expander device responds with SMP ZONE MANAGEMENT CLIENT VIOLATION, and the SAS address of the registered zone management client
3. HBA_2 issues a DISCOVER request to the SAS address of the registered zone management client
 - a. If the DISCOVER responds and the registered zone management client is active, then an error condition shall be reported, and the system policy will determine whether the revision can be made (see b).
 - b. If the DISCOVER does not respond then the registered zone management client may issue the SMP ZONE MANAGEMENT UPDATE request to each zoning expander device and change the registered zone management client to the address of HBA_2.
4. HBA_2 issues the SMP CONFIGURE ZONE PERMISSION to each zoning expander device in the ZPSDS.

.... The update proceeds as in case 1 (responses to HBA_2).

Case 3: HBA_1 and HBA_2 concurrently issue an SMP CONFIGURE ZONE PERMISSION request.

PROXY DEVICE (defined in 06-215):

1. HBA_2 issues SMP CONFIGURE ZONE PERMISSION request to the proxy device.
2. HBA_1 issues SMP CONFIGURE ZONE PERMISSION request to the proxy device.

Assumption: HBA_2 request arrives at the zone proxy device first.

3. The proxy device responds to HBA_2 with the FUNCTION RESULT field set to ZONE CONFIGURATION ACKNOWLEDGE, indicating that the proxy device successfully received the request and is scheduled to apply the change to all zoning expander devices.
4. The proxy device responds to HBA_1 with the FUNCTION RESULT field set to ZONE PROXY DEVICE BUSY, indicating that the proxy device has already accepted and is processing a prior SMP CONFIGURE ZONE PERMISSION request. HBA_1 may retry or abort determined by system policy.

.... The update proceeds as in case 1 (responses to HBA_2).

REGISTERED ZONE MANAGEMENT CLIENT (defined in 06-286):

Assumptions: HBA_1 is the REGISTERED ZONE MANAGEMENT CLIENT.

1. HBA_2 issues the SMP CONFIGURE ZONE PERMISSION to each zoning expander device in the ZPSDS.
2. HBA_1 issues the SMP CONFIGURE ZONE PERMISSION to each zoning expander device in the ZPSDS.
3. Each zoning expander device responds to HBA_2 with SMP ZONE MANAGEMENT CLIENT VIOLATION, and the SAS address of the registered zone management client
4. Each zoning expander device responds to HBA_1 with SMP FUNCTION ACCEPTED or an error condition. Any error condition may result in a retry by the REGISTERED ZONE MANAGEMENT CLIENT
5. HBA_2 issues a DISCOVER request to the SAS address of the registered zone management client
6. When the DISCOVER responds and the registered zone management client is active, then an error condition shall be reported, and the system policy will determine whether the revision can be made.
7. Each zoning expander device responds to HBA_1 with SMP FUNCTION ACCEPTED or an error condition. Any error condition may result in a retry by the REGISTERED ZONE MANAGEMENT CLIENT

.... The update proceeds as in case 1.

Case 4: Failure modes for discussion.

Detailed failure modes are not detailed here, due to lack of time prior to the next face to face meeting, and will be added in the next revision of this document.

Key failures for discussion:

PROXY DEVICE

- Proxy device fails – failover and re-election.
- HBA device fails during proxy process of SMP CONFIGURE ZONE PERMISSION to each zoning expander device in the ZPSDS.
- A zoning expander device fails to respond during SMP CONFIGURE ZONE PERMISSION to each zoning expander device in the ZPSDS.

REGISTERED ZONE MANAGEMENT CLIENT

- HBA device fails during MP CONFIGURE ZONE PERMISSION to each zoning expander device in the ZPSDS.
- A zoning expander device fails to respond during SMP CONFIGURE ZONE PERMISSION to each zoning expander device in the ZPSDS.

Actions for the discussion group

1. Benefit of shadow permission table to ensure ZPSDS is always consistent.
2. Registered management device vs. proxy device for zone management.
3. Are most cases addressed by these usage models
4. Error cases