

To: T10 Technical Committee
 From: Rob Elliott, HP (elliott@hp.com)
 Date: 19 June 2006
 Subject: 06-189r2 SAS-2 Allow table-to-table expander attachment

Revision history

Revision 0 (18 April 2006) First revision

Revision 1 (26 April 2006) Incorporated comments from 20 April 2006 SAS-2 zoning WG - got rid of edge expander/fanout expander designations and edge expander device sets altogether.

Revision 2 (19 June 2006) Incorporated comments from May 2006 SAS protocol WG.

Related documents

sas2r04 - Serial Attached SCSI - 2 (SAS-2) revision 4

06-078 - SAS-2 Expander routing table (Steve Johnson, LSI Logic)

06-097 - SAS-2 Discover Configuring bit (Tim Symons, PMC-Sierra)

Overview

As SAS-1 was originally defined, a constant refrain was to keep expanders as simple (and "cheap") as possible. The topology limited to a single *fanout expander* (a more complex device with routing tables) attached to up to 128 *edge expanders* (a simpler device without routing tables). The edge expanders use *subtractive routing* (named after the PCI bus "subtractive decode" terminology) to forward connection requests to the fanout expander when they do not match any attached devices.

Towards the end of SAS-1 development, however, it became apparent that expanders could implement small routing tables without much difficulty and expense, so limited route table capability was defined for edge expanders. Edge expanders were grouped into edge expander device sets (by proposal 02-359r3), and the fanout expander could attach to one edge expander set per phy. An edge expander routing table contains no more than 128 SAS addresses, covering the entire edge expander device set. A fanout expander routing table could contain $128 * (\text{number of phys})$ SAS addresses, the largest being $128 * 128 = 16,384$.

As SAS-2 is being defined, however, edge expanders are implementing larger routing tables, ignoring the limit of 128 SAS addresses per edge expander device set. Zoning expander devices may effectively contain a routing table for their subtractive ports if the zone group information is stored for every SAS address behind the subtractive port. There is fundamentally no difference between an edge expander and a fanout expander except what they claim to be in their IDENTIFY address frames, and there are few implementations of fanout expanders.

SAS-2 should include several changes to reflect the evolving market:

- a) Eliminate edge expander device sets
 - A) Allow more than 128 SAS addresses in the routing table for any given phy
 - B) This removes the limit of 16,384 SAS addresses per SAS domain
- b) Eliminate edge expander and fanout expander concepts
 - A) Just have plain "expander"s
 - B) Allow fanout (or edge) expanders to be attached to each other (i.e., remove the limit of one fanout expander per SAS domain)
- c) Allow table-to-table attachments for self-configuring expander devices
 - A) Continue to prohibit this for externally configurable expander devices, so the SAS-1.1 discovery and configuration algorithms are unaffected
 - B) Continue to allow self-configuring expander devices that do not support table-to-table
 - C) On Mini SAS 4x connectors leading to table routing phys supporting table attachment, use the permissive key 4 (like an end device)
 - D) On external SAS connectors leading to table routing phys supporting table attachment, use both diamond and circle icons (like an end device)
 - E) This removes the limit of 16,384 SAS addresses per SAS domain
- d) Add a TABLE TO TABLE SUPPORTED bit in byte 10 of the SMP REPORT GENERAL response indicating that the self-configuring expander device allows table-to-table attachments
 - A) Not supported by externally configurable expander devices
- e) Continue to prohibit loops

- f) Separate the discover process (performed by all SMP initiator ports, including those in self-configuring expander devices, to determine what SAS addresses are accessible) from the configuration subprocess (performed by any SMP initiator port, to fill in the routing tables in any externally configurable expander devices)
 - A) Discover process honors SAS-2 table-to-table attachment if detected in an expander device with the TABLE TO TABLE SUPPORTED bit set to one
 - B) Configuration process still applies the SAS-1/SAS-1.1 algorithms, and skips over any table-to-table attachments it detects
 - C) Future versions of SAS may consider externally configurable expanders and the configuration subprocess as obsolete.

Figure 1 shows a basic example.

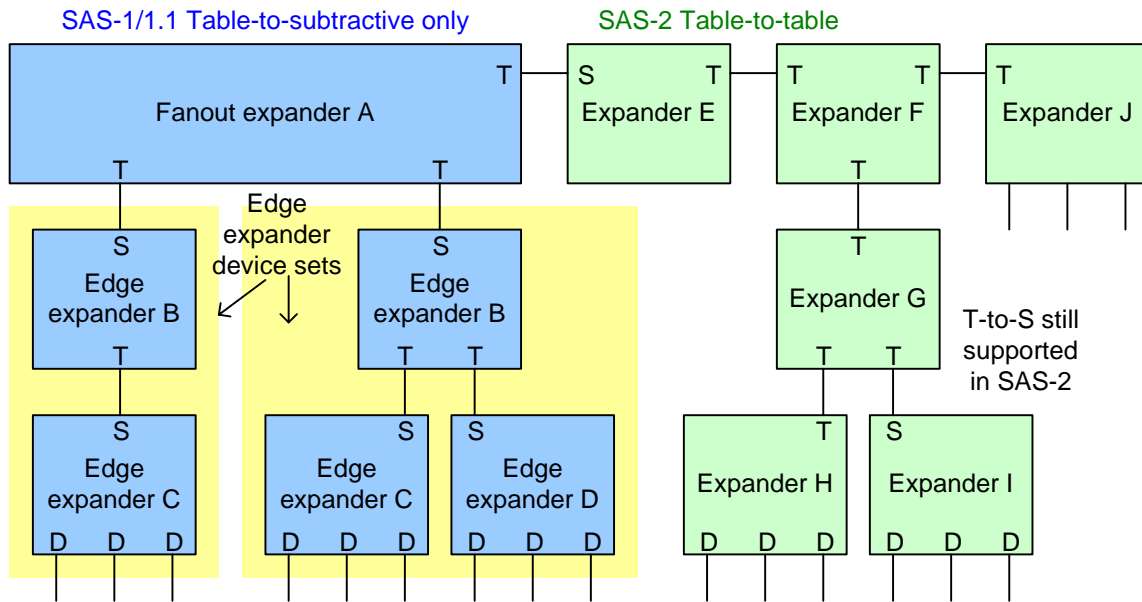


Figure 1 — Table-to-table example

A SAS domain could contain any mix of:

- a) externally configurable expander device - route table filled in with the SAS-1/SAS-1.1 non-optimized configuration process
- b) externally configurable expander device - route table filled in with the SAS-1.1 optimized configuration process
- c) self-configuring expander device not supporting table-to-table
 - A) using the SAS-1.1 discover process
 - B) using the SAS-2 discover process (tolerant of table-to-table)
- d) self-configuring expander device supporting table-to-table
- e) SMP initiator ports using the SAS-1.1 discover and configuration process. These skip over table-to-table attachments, which is desired for configuration purposes but limiting for discover purposes (if a self-configuring expander device supports table-to-table, this initiator port won't understand it can talk to the SAS addresses beyond that point)
- f) SMP initiator ports using the SAS-2 discover and configuration process.

Additional changes incorporated:

- a) add 4.1.x application client and device server section. Management application client and management device server terms are used more by this proposal, but are not formally introduced.
- b) upgrade ATA/ATAPI-7 V1 references to ATA8-AAM or ATA8-ACS as appropriate

Suggested changes to chapter 3

2.3 References under development

...

~~ISO/IEC 14776-971, AT Attachment with Packet Interface-7 Volume 1 (ATA/ATAPI-7 V1) (ANSI INCITS-397-2005)~~

ISO/IEC 14776-971, AT Attachment with Packet Interface-7 Volume 3 (ATA/ATAPI-7 V3)(Serial ATA) (ANSI INCITS 397-2005)

ISO/IEC 14776-861, ATA Attachment-8 Architecture Model (ATA8-AAM) (T13/1700-D)

ISO/IEC 14776-871, ATA Attachment-8 Architecture Model (ATA8-ACS) (T13/1699-D)

...

3.1.6 application client: An object that is the source of SCSI commands [and task management function requests](#) (see SAM-3), ATA commands (see ~~ATA/ATAPI-7 V1~~[ATA8-AAM](#)), or management function requests. See [4.1.x](#).

3.1.10 AT Attachment (ATA): A standard for the internal attachment of storage devices to hosts. See ~~ATA/ATAPI-7 V1~~[ATA8-AAM](#).

3.1.11 ATA device: A storage peripheral (analogous to a SCSI target device). See ~~ATA/ATAPI-7 V1~~[ATA8-AAM](#).

3.1.13 ATA host: A host device that originates requests to be processed by an ATA device (analogous to a SCSI initiator device). See ~~ATA/ATAPI-7 V1~~[ATA8-AAM](#).

3.1.28 externally configurable expander device: An [non-self-configuring](#) expander device that contains an expander route table ~~that is configured with expander route entries~~ [configurable by the SMP CONFIGURE ROUTE INFORMATION function](#) (see 10.4.3.12). See 4.1.5.

3.1.xx configuration subprocess: A subprocess invoked from the discover process to configure an [externally configurable expander device](#). See 4.8.

3.1.46 device server: An object ~~within a SAS target device~~ that processes SCSI tasks (see SAM-3), ATA commands (see ~~ATA/ATAPI-7 V1~~[ATA8-AAM](#)), or management functions. See [4.1.x](#).

3.1.xx discover process: The process performed by a management application client to discover all the SAS devices and expander devices in the SAS domain, invoking the configuration subprocess as needed. See 4.7.

~~**3.1.56 edge expander device:** An expander device that is part of a single edge expander device set (see 3.1.57). See 4.1.5.~~

~~**3.1.57 edge expander device set:** A group of one or more edge expander devices (see 3.1.56) that may be attached to no more than one other edge expander device set or one fanout expander device (see 3.1.75). See 4.1.8.2.~~

3.1.59 enclosure in port: A set of expander phys with subtractive routing attributes using the same external connector (see 5.2.3.3). See 4.6.2.

3.1.60 enclosure out port: A set of expander phys with table routing attributes [in an expander device that does not support table-to-table attachment](#) using the same external connector (see 5.2.3.3). See 4.6.2.

3.1.xx enclosure universal port: A set of expander phys with table routing attributes [in an expander device that supports table-to-table attachment using the same external connector](#) (see 5.2.3.3). See 4.6.2.

3.1.65 expander device: A device that is part of the service delivery subsystem and facilitates communication between SAS devices. See 4.1.5.

3.1.69 expander route entry: A SAS address and an enable/disable bit in an expander route table (see 4.6.7.3).

3.1.70 expander route index: A value used in combination with a phy identifier to select an expander route entry in an expander route table [in an externally configurable expander device](#) (see 4.6.7.3).

3.1.71 expander route table: A table of expander route entries within an expander device. The table is used by the expander function to resolve connection requests. See 4.6.7.3.

3.1.xx fanout cable assembly: [A cable assembly that attaches a wide connector on one end to multiple narrow connectors on the other end. See 5.2.4.1.3.](#)

~~**3.1.75 fanout expander device:** An expander device that is capable of being attached to two or more edge expander device sets. See 4.1.5.~~

3.1.155 read data: Data transferred to the [SCSI](#) application client's data-in buffer from the [SCSI](#) device server, as requested by the Send Data-In transport protocol service (see 10.2.1.6).

3.1.xx route table optimization: [Configuration subprocess algorithm that reduces the number of entries required in an expander route table in an externally configurable expander device. See 4.7.3.](#)

3.1.192 SCSI initiator device: A SCSI device containing [SCSI](#) application clients and SCSI initiator ports that originates device service and task management requests to be processed by a SCSI target device and receives device service and task management responses from SCSI target devices. See SAM-3.

3.1.193 SCSI initiator port: A SCSI initiator device object that acts as the connection between [SCSI](#) application clients and the service delivery subsystem through which indications and responses are routed. See SAM-3.

3.1.195 SCSI target device: A SCSI device containing logical units and SCSI target ports that receives device service and task management requests for processing and sends device service and task management responses to SCSI initiator devices. See SAM-3.

3.1.196 SCSI target port: A SCSI target device object that contains a task router and acts as the connection between [SCSI](#) device servers and task managers and the service delivery subsystem through which requests and confirmations are routed. See SAM-3.

3.1.xx self-configuring expander device: An expander device containing an SMP initiator port and a management application client to perform the discover process (see 4.7) and to configure its own expander route table. See 4.7.5.

3.1.236 subtractive routing attribute: The attribute of an [edge](#) expander phy that indicates it may be used by the expander connection manager to route connection requests not resolved using the direct routing method or table routing method. See 4.6.7.1.

3.1.237 subtractive routing method: The method the expander connection manager uses to route connection requests [to an attached expander device](#) not resolved using the direct routing method or table routing method to an expander device. See 4.6.7.1.

3.1.239 table routing attribute: The attribute of an expander phy that indicates it may be used by the expander connection manager to route connection requests using an expander route table. See 4.6.7.1.

3.1.240 table routing method: The method the expander connection manager uses to route connection requests to an [attached](#) expander device using an expander route table. See 4.6.7.1.

3.1.244 task management function: A task manager service capable of being requested by an application client to affect the processing of one or more tasks. See SAM-3.

~~**3.1.245 task manager:** An [agent within the device server](#) object within a SAS target device that [controls the sequencing of SCSI commands and](#) processes [SCSI](#) task management functions. See SAM-3.~~

3.1.264 write data: Data transferred from the [SCSI](#) application client's data-out buffer to the [SCSI](#) device server, as requested by the Request Data-Out transport protocol service (see 10.2.1.8).

Suggested changes to chapter 4 (Architecture)

4.1 Architecture

4.1.1 Architecture overview

A SAS domain (see 4.1.7) contains one or more SAS devices and a service delivery subsystem. A SAS domain may be a SCSI domain (see SAM-3).

A SAS device (see 4.1.4) contains one or more SAS ports (see 4.1.3). A SAS device may be a SCSI device (see SAM-3).

A SAS port (see 4.1.3) contains one or more phys (see 4.1.2). A SAS port may be a SCSI port (see SAM-3).

The service delivery subsystem (see 4.1.6) in a SAS domain may contain expander devices (see 4.1.5).

Expander devices contain expander ports (see 4.1.3) and one SMP port.

An expander port contains one or more phys (see 4.1.2).

An expander device shares its phys with the SAS device(s) contained within the expander device.

Figure 2 shows the class diagram for a SAS domain, showing the relationships between SAS domain, SCSI domain, service delivery subsystem, expander device, expander port, SAS device, SCSI device, SAS port, SCSI port, and phy classes. Not all attributes are shown.

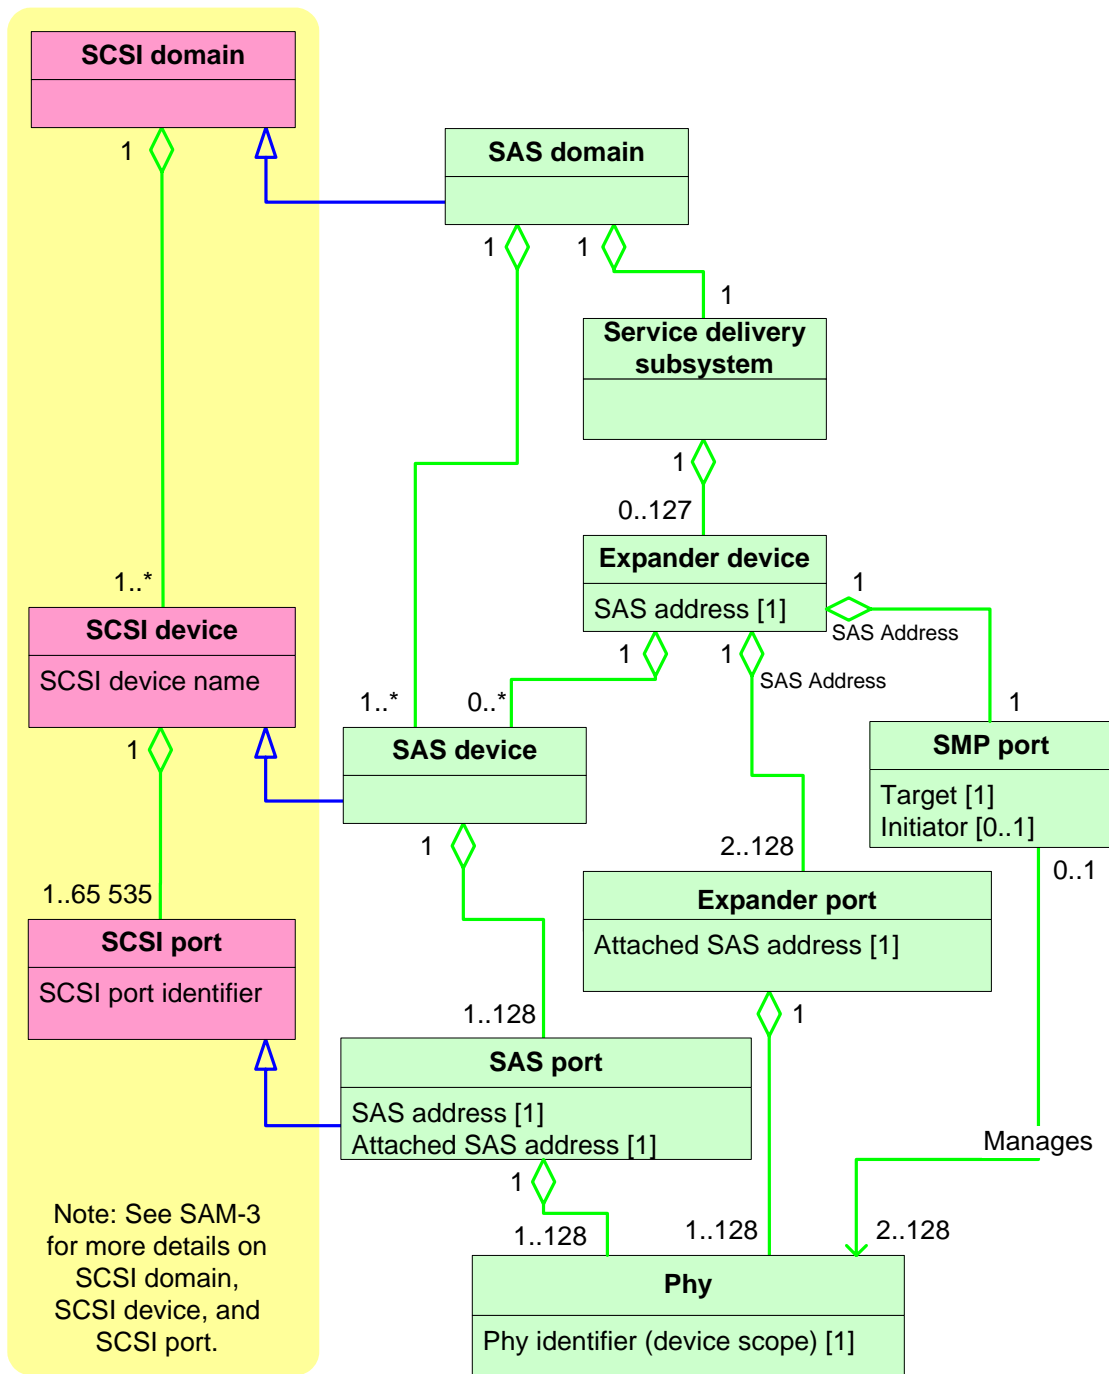


Figure 2 — SAS domain class diagram

Editor’s Note 1: Deferred to another proposal: add application client and device server to this class diagram, and bring it up to date with SAM-4. In SAM-4, SCSI device contains SCSI target device, which contains level 1 hierarchical logical unit, which contains logical unit, which contains device server. SCSI device also contains SCSI initiator device, which contains application client. All the SAS UML diagrams should switch to the “contains” method rather than the inheritance method.

4.1.2 Physical links and phys

...

During the identification sequence (see 7.9), a phy:

- a) transmits an IDENTIFY address frame including the device type (i.e., end device, ~~edge expander device~~, or ~~fanout~~ expander device) of the device containing the phy, the SAS address of the SAS port or expander device containing the phy, the device name of the SAS device or expander device containing the phy, the phy identifier, and bits specifying the SSP initiator phy capability, STP initiator phy capability, SMP initiator phy capability, SSP target phy capability, STP target phy capability, and SMP target phy capability.
- b) receives an IDENTIFY address frame containing the same set of information from the attached phy, including the attached device type, the attached SAS address, the attached device name, the attached phy identifier, and bits indicating the attached SSP initiator phy capability, attached STP initiator phy capability, attached SMP initiator phy capability, attached SSP target phy capability, attached STP target phy capability, and attached SMP target phy capability.

...

4.1.x Application clients and device servers

This standard defines the following application clients:

- a) a SCSI application client (see SAM-3) is the source of SCSI commands and task management function requests. A SCSI application client uses an SSP initiator port to interface to the service delivery subsystem;
- b) an ATA application client (see ATA8-AAM) is the source of ATA commands and device management operation requests behind an STP initiator port. An ATA application client uses an STP initiator port to interface to the service delivery subsystem; and
- c) a management application client is the source of management function requests behind an SMP initiator port. A management application client uses an SMP initiator port to interface to the service delivery subsystem.

This standard defines the following device servers:

- a) a SCSI device server (see SAM-3) processes SCSI commands. A SCSI device server uses an SSP target port to interface to the service delivery subsystem;
- b) an ATA device server (see ATA8-AAM) processes ATA commands. An device server uses an STP target port to interface to the service delivery subsystem; and
- c) a management device server processes management functions. A management device server uses an SMP target port to interface to the service delivery subsystem;

4.1.5 Expander devices (~~edge expander devices and fanout expander devices~~)

Expander devices are part of the service delivery subsystem and facilitate communication between multiple SAS devices. Expander devices contain two or more external expander ports. Each expander device contains one SMP target port and one management device server, contains one SMP initiator port and one management application client if it is self-configuring and may contain one SMP initiator port and one management application client if it is not self-configuring, and may contain SAS devices (e.g., an expander device may include an SSP target port for access to a logical unit with a peripheral device type set to 0Dh (i.e., enclosure services device) (see SPC-3 and SES-2)).

Figure 3 shows an expander device.

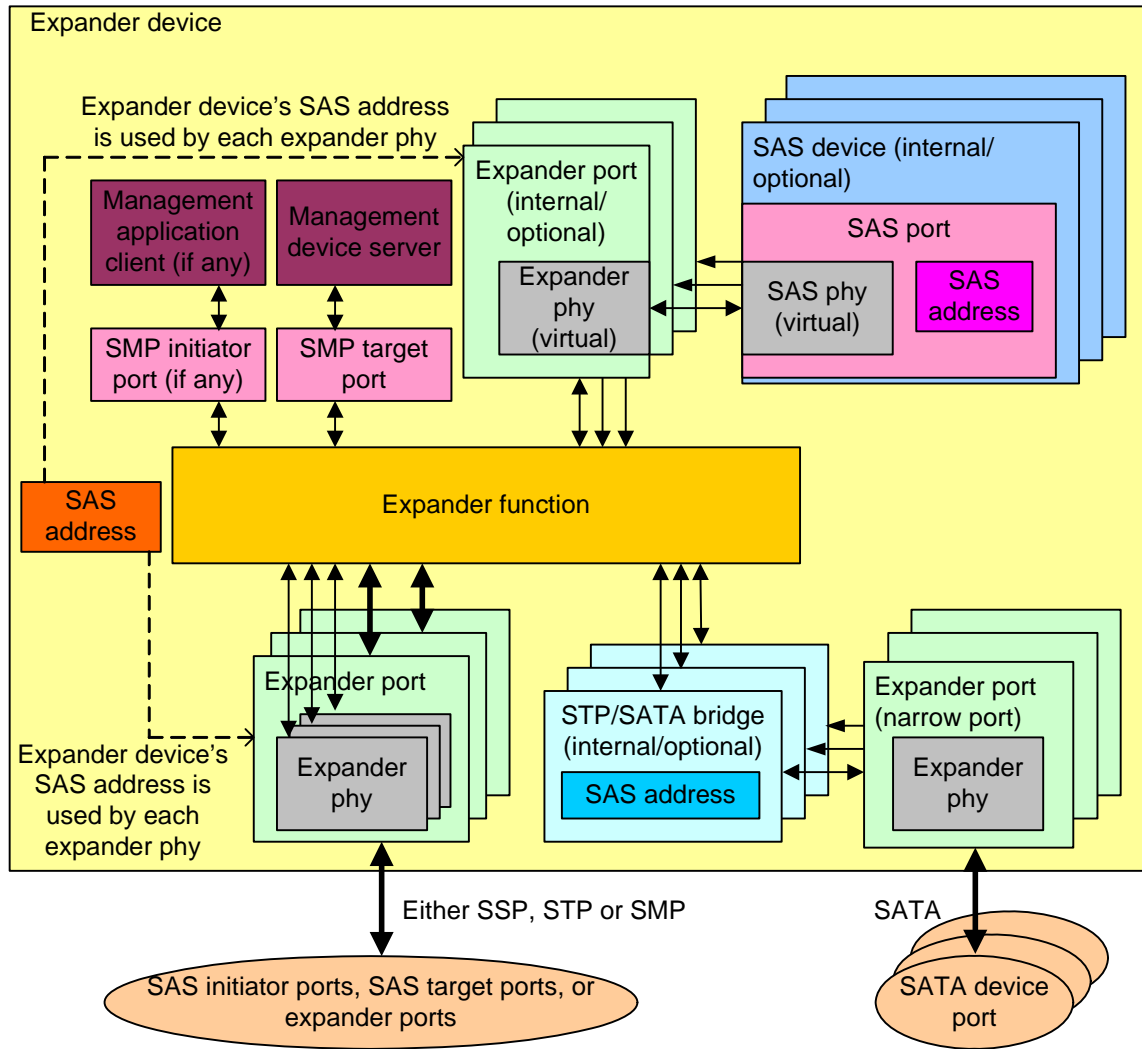


Figure 3 — Expander device

There are two types of expander devices:

- a) edge-expander devices; and
- b) fanout-expander devices.

An edge-expander device is part of an edge-expander device set (see 4.1.8.2). A fanout expander device is an expander device capable of being attached to two or more edge-expander device sets.

See 4.6 for a detailed model of an expander device.

Figure 4 shows/defines the expander device classes, showing the relationship between the following classes:

- a) expander device;
- b) edge-expander device;
- e) edge-expander device set; and
- d) fanout-expander device.




Figure 4 — Expander device class diagram [\[modified to eliminate edge, fanout, and set\]](#)

Each expander phy has one of the following routing attributes (see 4.6.7.1):

- a) direct routing attribute;
- b) table routing attribute; or
- c) subtractive routing attribute.

~~Edge expander devices may contain phys with the subtractive routing attribute. A fanout expander device shall not contain any phys with the subtractive routing attribute. Edge expander devices and fanout expander devices may also contain phys with direct routing and table routing attributes.~~

Expander devices with expander phys with the table routing attribute contain an expander route table. ~~The expander route table may be configurable.~~ An [externally configurable](#) expander device ~~with a configurable route table~~ depends on a management application client within the SAS domain to use the discover process (see 4.7) ~~and the configuration subprocess (see 4.8)~~ to configure the expander route table. ~~An expander device with expander phys with the table routing attribute that does not have a configurable route table shall be self-configuring, and shall~~ [A self-configuring expander device](#) contains a management application client and SMP initiator port to perform the discover process ~~(see 4.7)~~ to configure its own expander route table.

4.1.8 Expander device topologies

~~4.1.8.1 Expander device topology overview~~

More than one expander device may be part of a service delivery subsystem.

~~Expander devices shall be attached in a tree topology with no loops (i.e., there shall be no more than one path from any end device or expander device to any other end device or expander device in the service delivery subsystem).~~

~~Expander device phys with table routing attributes shall only be attached to expander device phys with table routing attributes if both expander devices support table-to-table attachment (i.e., they each have the TABLE TO TABLE SUPPORTED bit set to one in the REPORT GENERAL response (see 10.4.3.3)).~~

~~4.1.8.2 Edge expander device set~~

~~One or more edge expander devices may be grouped into an edge expander device set. The number of edge expander devices and the routing attributes (see 4.6.7.1) of the expander phys in the edge expander devices within an edge expander device set shall be established when the edge expander device set is constructed. The method used to construct edge expander sets is outside the scope of this standard. The routing method used by an expander phy depends on the type of device to which it is attached (see 4.6.7.1).~~

~~An edge expander device set is bounded by expander phys with the direct routing attribute that are either:~~

- ~~a) attached to end devices; or~~
- ~~b) not attached to any device,~~

~~and by expander phys with the subtractive routing attribute that are:~~

- ~~a) attached to expander phys of a fanout device;~~
- ~~b) attached to expander phys with subtractive routing attributes on another edge expander device set;~~
- ~~c) attached to an end device; or~~
- ~~d) not attached to any device.~~

Phys with table routing attributes within an edge expander device set are used to provide attachments between edge expander devices in the edge expander device set.

The number of SAS addresses used by an edge expander device set shall not exceed 128. This includes SAS addresses for:

- a) the edge expander devices within the edge expander device set;
- b) SAS devices contained in the expander devices;
- c) STP target ports in STP/SATA bridges contained in the expander devices; and
- d) the end devices potentially attached to the edge expander device set.

To avoid an overflow of an edge expander route index during the discover process configuration subprocess (see 4.7)(see 4.8), an edge expander device set SAS domain containing an externally configurable expander device shall be constructed such that the number of edge expander route indexes available for each edge expander phy with the table routing attribute attached to another expander device is greater than or equal to the number of SAS addresses addressable through that edge expander phy.

Figure 5 shows an edge expander device set.




Figure 5 — Edge expander device set [delete this figure]

4.1.8.3 Expander device topologies

An edge expander device set shall not be attached to more than one fanout expander device.

An edge-expander device set may be attached to one other edge-expander device set if:

- a) there are only two edge-expander device sets in the SAS domain;
- b) the edge-expander device sets are attached using expander phys with subtractive-routing attributes; and
- c) there are no fanout expander devices in the SAS domain.

No more than one fanout expander device shall be included in a SAS domain. The fanout expander device may be attached to up to 128 edge-expander device sets or SAS ports.

Figure 6 shows a SAS domain with the maximum number of expander device sets.

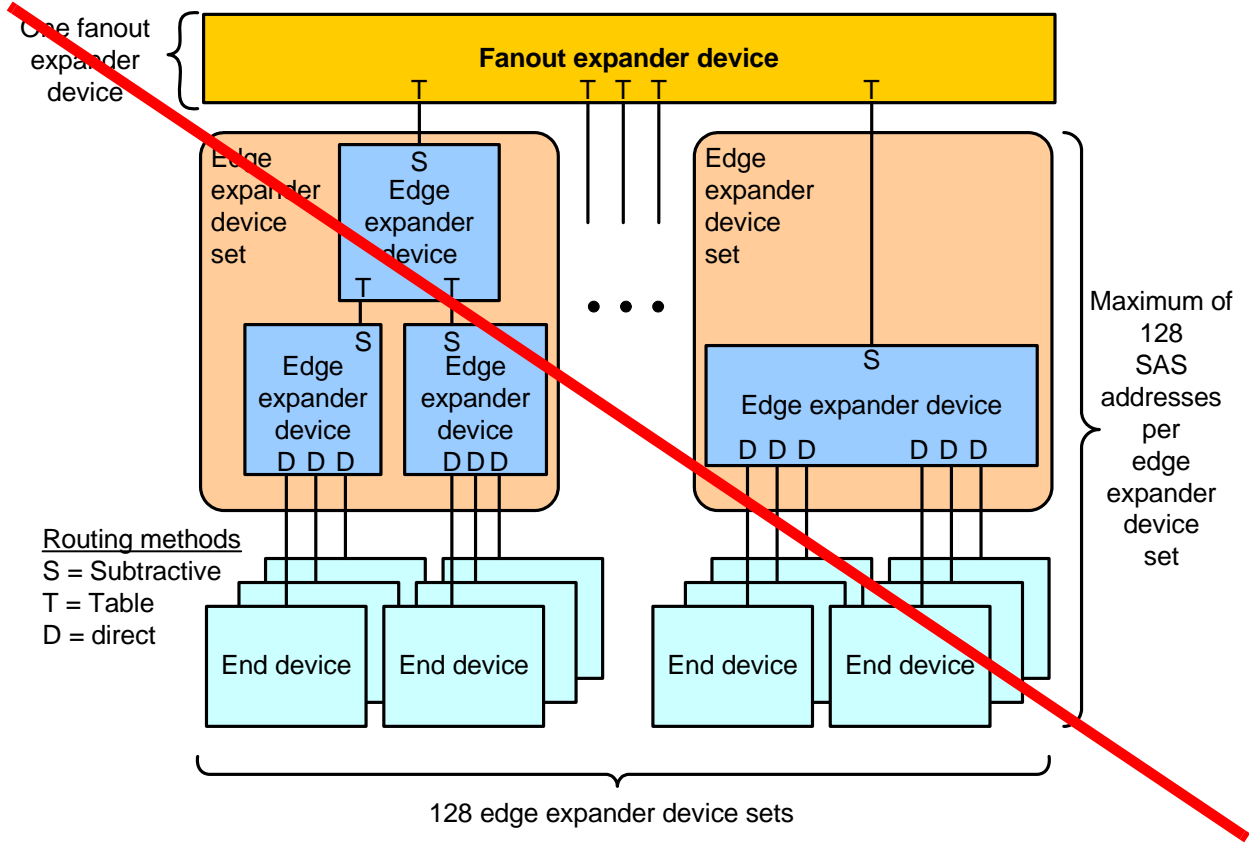


Figure 6 — ~~Maximum expander device set topology~~ [\[delete this figure\]](#)

End devices may be attached directly to the fanout expander device, as shown in figure 7.

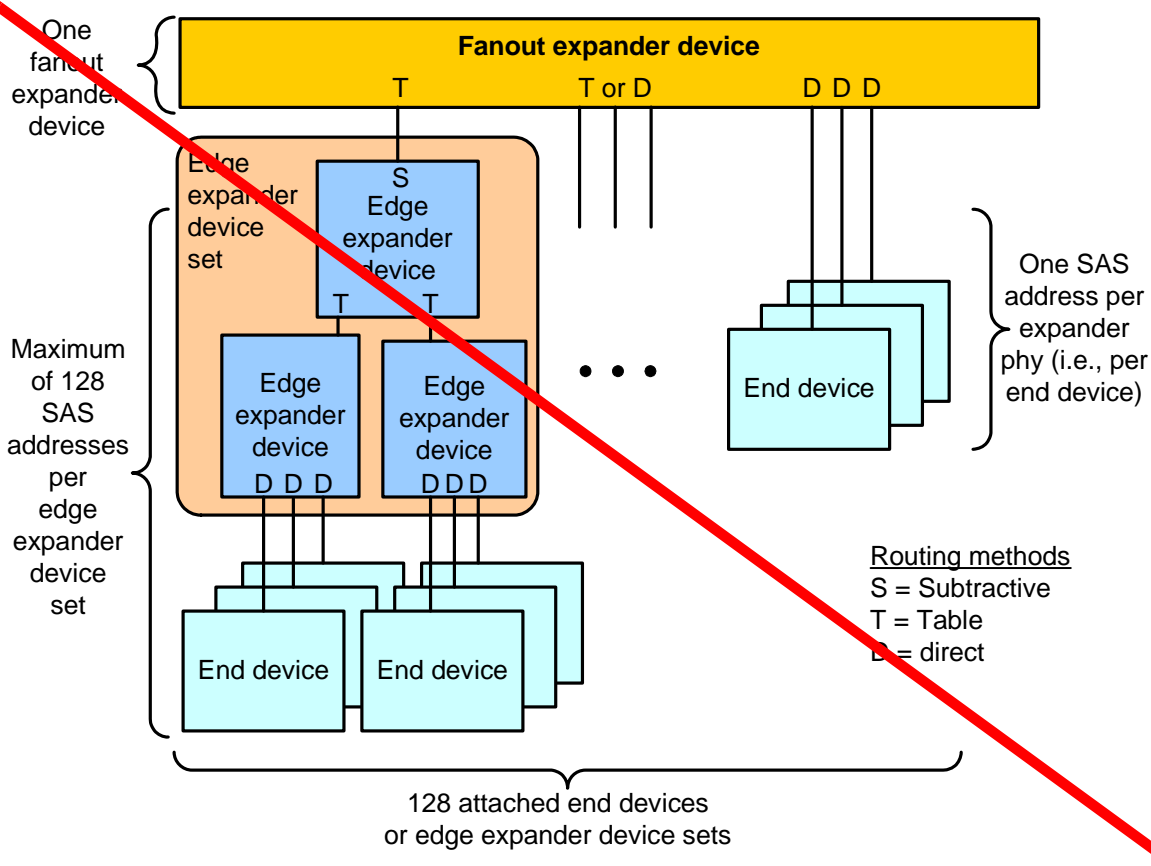


Figure 7 — Fanout expander device topology [delete this figure]

SAS initiator devices and SAS target devices may be attached to any of the edge expander devices within an edge expander device set and at most, two edge expander device sets may be attached together without a fanout expander, as shown in figure 8.

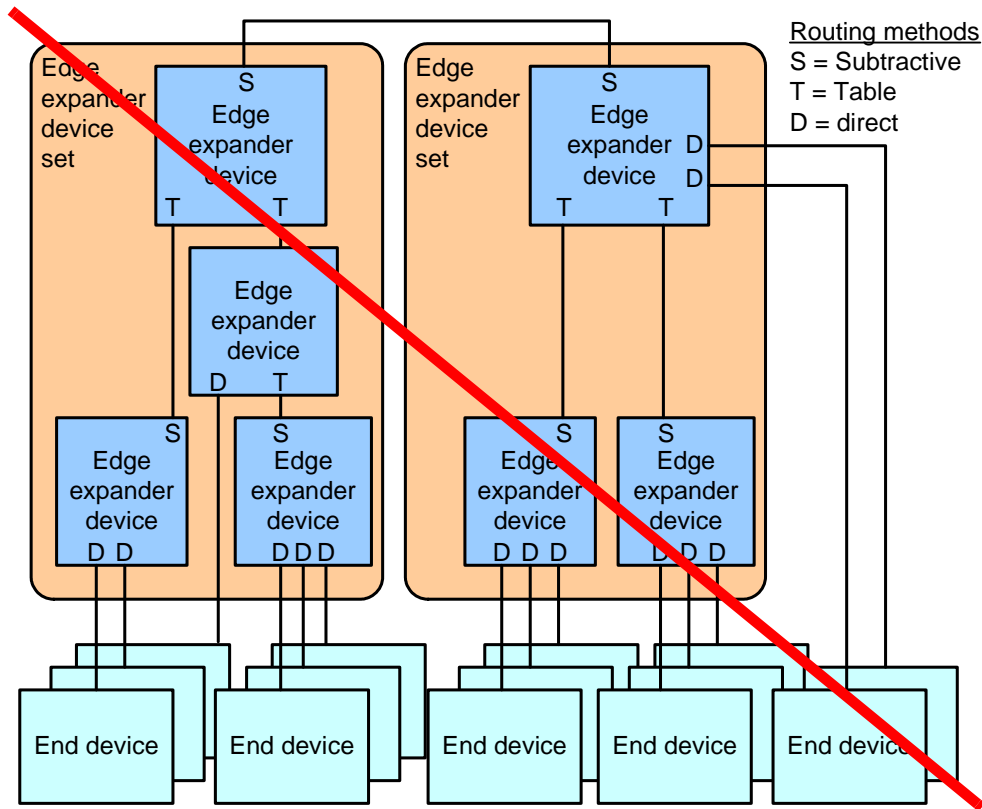


Figure 8 — Edge expander device set to edge expander device set topology [\[delete this figure\]](#)

Figure 9 shows an example of an expander topology with one expander device.

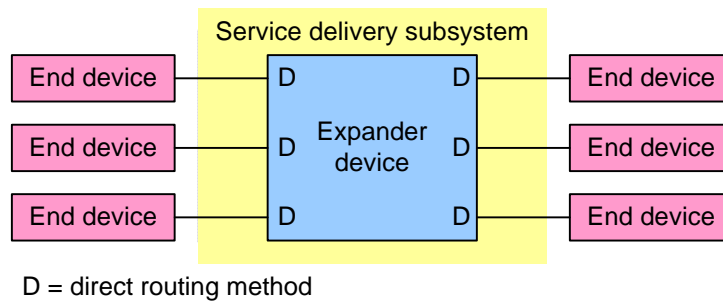


Figure 9 — Single expander device topology example

Figure 10 shows examples of expander topologies with multiple expander devices.

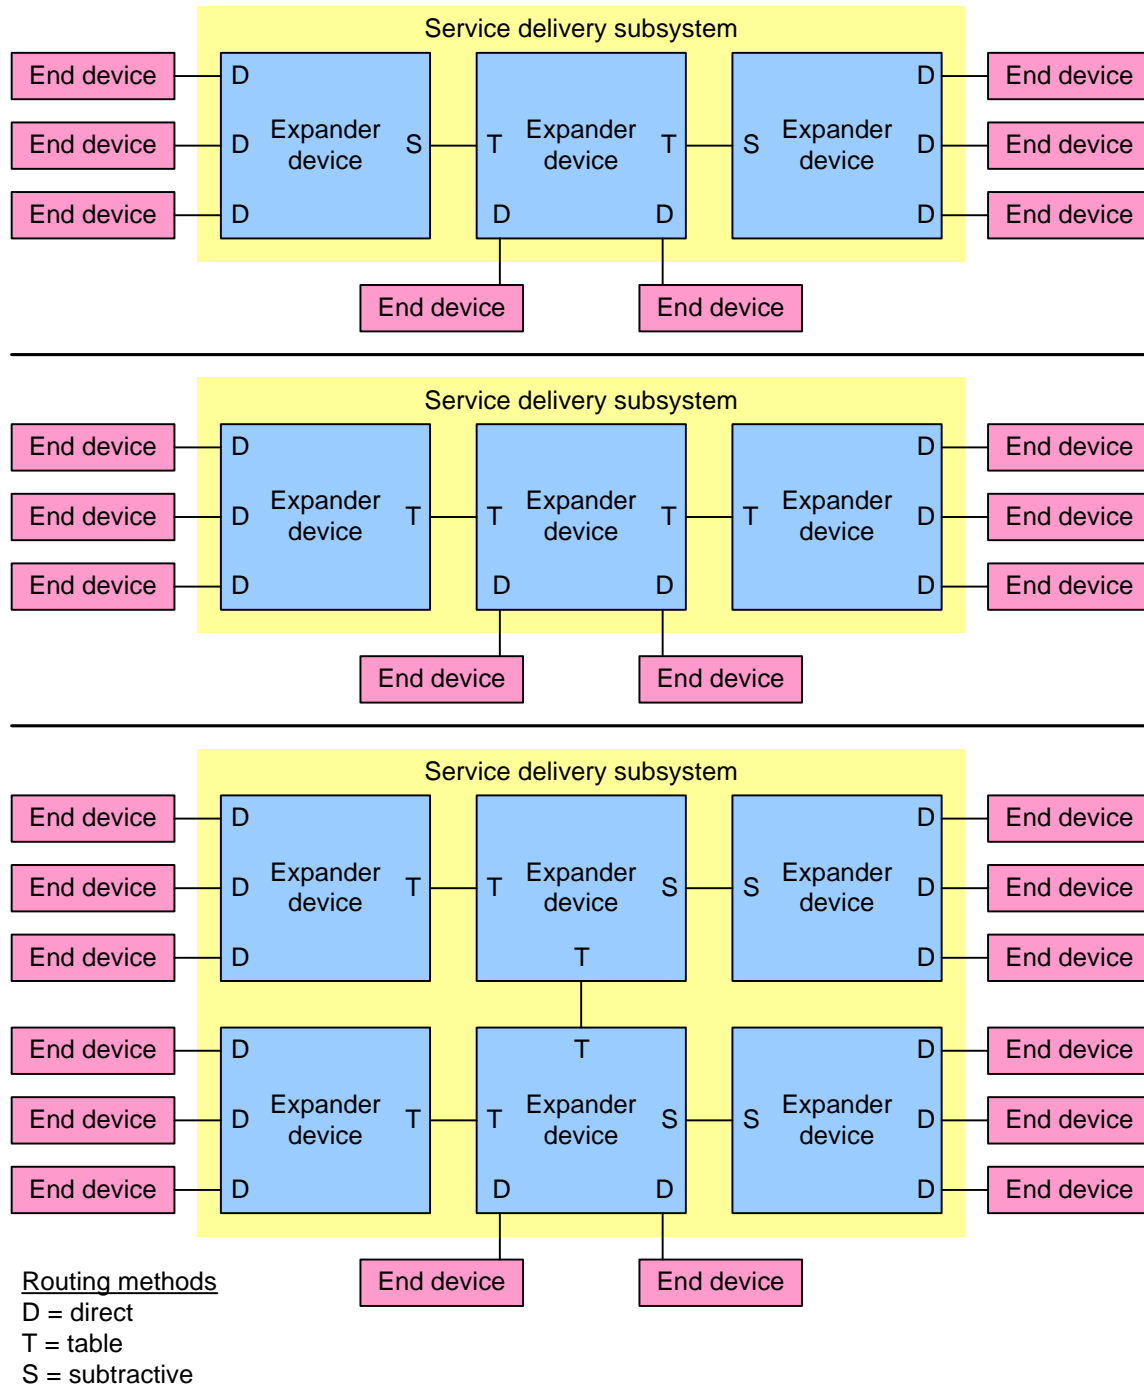


Figure 10 — Multiple expander device topologies

4.4 Resets

4.4.2.2 Additional hard reset processing by SAS ports

...

If the port processing the hard reset is an STP port that is not in an STP/SATA bridge, the STP target device shall perform the actions defined for power-on or hardware reset in [ATA/ATAPI-7-V1](#) and [ATA8-AAM](#).

4.5 I_T nexus loss

When a SAS port receives OPEN_REJECT (NO DESTINATION), OPEN_REJECT (PATHWAY BLOCKED), or an open connection timeout occurs in response to a connection request, it shall retry the connection request until:

- a) the connection is established;
- b) for SSP target ports, the time indicated by the I_T NEXUS LOSS field in the Protocol-Specific Port mode page expires; or
- c) the I_T nexus loss timer, if any, expires (see [4.7.1](#), 8.2.2.1, 10.2.7.2, and 10.4.3.11).

I_T nexus loss is handled by the port layer state machines (see 8.2.2.3).

If the I_T nexus loss timer expires in an SSP port, then the port shall send a Nexus Loss event notification to the SCSI application layer (see 10.2.5) and the SCSI device shall perform the actions defined for I_T nexus loss in SAM-3.

If the I_T nexus loss timer expires in an STP target port, then the port shall abort any commands for the lost STP initiator port. If the STP target port is in an STP/SATA bridge, the STP/SATA bridge shall originate a link reset sequence to the SATA device.

If the I_T nexus loss timer expires in an STP initiator port, then the port shall consider any commands for the lost STP target port to be completed with an error. This serves as a nexus loss event (see ATA8-AAM).

If the I_T nexus loss timer expires in an SMP initiator port, then the port shall stop attempting to establish connections to the lost SMP target port.

If the I_T nexus loss timer expires in an initiator port, then a management application client should cause a link reset sequence on the phy(s) attached to the lost target port (e.g, if directly attached, the phys in the initiator port; if attached via expander device(s), the phys in the expander device closest to and attached to the target port).

4.6 Expander device model

4.6.2 Expander ports

An external expander port contains one or more physical phys (see 4.1.2). Since each phy in the expander device has the same SAS address, expander ports are created based on the attached SAS addresses (see 4.1.3).

Each phy in an expander port shall have the same routing attribute (see 4.6.7.1), and the DISCOVER function (see 10.4.3.5) shall return the same value in the ROUTING ATTRIBUTE field for each phy in an expander port.

Each phy in an expander port containing phys with table routing attributes [in an externally configurable expander device](#) shall have the same number of routing table entries (see 4.6.7.3).

A set of expander phys with table routing attributes [in an expander device not supporting table-to-table attachments](#) using the same external connector (see 5.2.3.3) is called an enclosure out port. A set of expander phys with subtractive routing attributes using the same external connector is called an enclosure in port. [A set of expander phys with table routing attributes in an expander device supporting table-to-table attachments using the same external connector \(see 5.2.3.3\) is called an enclosure universal port.](#)

Each phy in an expander port shall have the same zone phy information (see 4.8.3.1). The zone phy information associated with each of the phys in an expander port is treated as the zoning properties of the expander port.

Each expander phy contains an expander link layer with an XL state machine (see 7.15) and an SL_IR state machine (see 7.9.5). The XL state machine in each expander phy within an expander port processes connection requests independently of the XL state machines in other expander phys.

An internal expander port contains a virtual phy with an expander link layer and a protocol-specific transport layer (e.g., to provide access as an SSP target port to a logical unit with a peripheral device type set to 0Dh (i.e., enclosure services device) (see SPC-3 and SES-2)).

Each expander device shall include one internal SMP port using the expander device's SAS address.

Any additional internal SAS ports shall be inside SAS devices contained in the expander device, and thus have SAS addresses different from that of the expander device. These SAS ports shall be attached to internal expander ports with virtual phys.

Each STP/SATA bridge shall have a unique SAS address. This SAS address is reported in the ATTACHED SAS ADDRESS field in the DISCOVER response (see 10.4.3.5) for the expander phy containing the STP/SATA bridge (i.e., the expander phy attached to the SATA device or SATA port selector).

4.6.5 Broadcast primitive processor (BPP)

The BPP receives Broadcasts from each expander phy or from the management device server on behalf of an expander phy and requests transmission of those Broadcasts on all expander ports except the expander port from which the Broadcast was received.

~~In a self-configuring expander device (see 4.1.5), the BPP requests transmission of a BROADCAST (CHANGE) when it completes configuration (see 10.4.3.3).~~

In a zoning expander device (see 4.8.2), Broadcasts are forwarded as described in 4.8.5.

4.6.7 Expander device routing

4.6.7.1 Routing attributes and routing methods

Each expander phy in an expander device shall support one of the following routing attributes:

- a) direct routing attribute;
- b) table routing attribute; or
- c) subtractive routing attribute.

The routing attributes allow the ECM to determine which routing method to use when routing connection requests to the expander phy:

- a) the table routing method routes connection requests to attached expander devices using an expander route table;
- b) the subtractive routing method routes unresolved connection requests to an attached expander device; or
- c) the direct routing method routes connection requests to attached end devices, the SMP port of an attached expander device, or SAS devices contained in the expander device.

Table 1 describes the routing methods that the ECM uses based on the routing attributes of a phy.

Table 1 — Routing attributes and routing methods

Routing attribute of a phy	Routing method used by ECM for the phy
Direct	Direct ^a
Table	Direct, if attached to an end device
	Direct, if attached to an expander device, for the SAS address of the expander device
	Table, if attached to an expander device, for SAS addresses beyond the expander device
Subtractive	Direct, if attached to an end device
	Subtractive, if attached to an expander device
^a If attached to an expander device, the ECM is only able to route to the expander device itself through a phy with the direct routing attribute	

An expander device may have zero or more phys with the table routing attribute.

An **edge** expander device shall have at most one defined port containing phys with the subtractive routing attribute. ~~Phys in a fanout expander device shall not have the subtractive routing attribute.~~

An edge expander device [that does not support table-to-table attachment](#) shall only use phys with the table routing attribute to attach to phys with the subtractive routing attribute in other edge expander devices within an edge expander device set.

If multiple phys within an expander device have subtractive routing attributes and are attached to expander devices, they shall attach to phys with identical SAS addresses (i.e., the same expander port).

If multiple phys within an expander device have subtractive routing attributes and are attached to expander devices that do not have identical SAS addresses, the application client that is performing the discover process (see 4.7) shall report an error in a vendor-specific manner.

4.6.7.2 Connection request routing

The ECM shall determine how to route a connection request from a source expander phy to a destination expander phy in a different expander port if the destination expander phy is enabled and operating at a valid physical link rate (e.g., the DISCOVER function reports a NEGOTIATED PHYSICAL LINK RATE field set to G1 (i.e., 8h) or G2 (i.e., 9h)) using the following precedence:

- 1) route to an expander phy with the direct routing attribute or table routing attribute when the destination SAS address matches the attached SAS address;
- 2) route to an expander phy with the table routing attribute when the destination SAS address matches an enabled SAS address in the expander route table;
- 3) route to an expander phy with the subtractive routing attribute; or
- 4) return an Arb Reject confirmation (see 4.6.6.3) to the source expander phy.

If the destination expander phy only matches an expander phy in the same expander port from which the connection request originated, then the ECM shall return an Arb Reject confirmation.

If the destination SAS address of a connection request matches a disabled SAS address in an expander route table, then the ECM shall ignore the match.

4.6.7.3 Expander route table

An expander device that supports the table routing method shall contain an expander route table. The expander route table is a structure that provides an association between destination SAS addresses and expander phy identifiers. Each association represents an expander route entry.

An [externally configurable](#) expander device reports the size of its expander route table ~~and indicates if the expander route table is configurable~~ in the SMP REPORT GENERAL response (see 10.4.3.3). Each expander route entry shall be disabled after power on.

A management application client may reference a specific expander route entry within an expander route table [in an externally configurable expander device](#) with the SMP REPORT ROUTE INFORMATION function (see 10.4.3.8) and the SMP CONFIGURE ROUTE INFORMATION function (see 10.4.3.13). [It may retrieve the contents of an expander route table in a self-configuring expander device with the SMP REPORT ROUTE TABLE function \(see 10.4.3.xx\).](#)

Editor's Note 2: 06-078 is proposing the REPORT ROUTE TABLE to return the contents of an expander routing table without following the phy/expander route index order in SAS-1/SAS-1.1.

Figure 11 shows a representation of an expander route table.

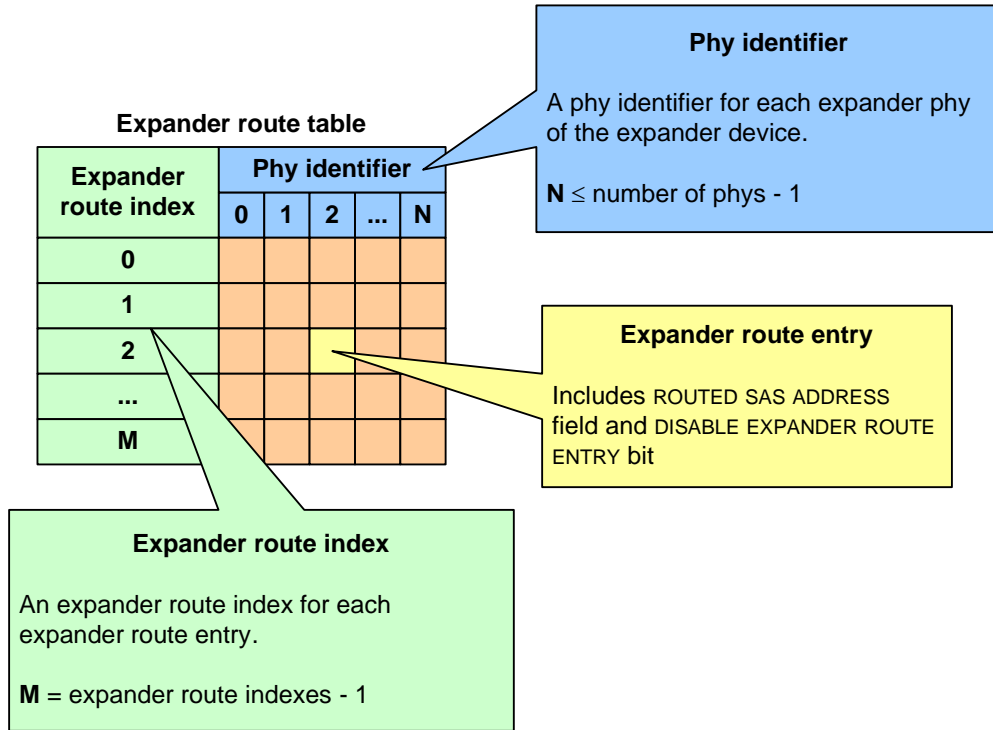


Figure 11 — Expander route table example

The number of end devices that may be attached to an edge expander device set is dependent on the number of expander route entries in the expander route table of the edge expander devices.

4.7 Discover process

4.7.1 Discover process overview

Management application clients direct an SMP initiator port to request SMP functions from an SMP target port. Management application clients are located in every SAS initiator device and every self-configuring expander device. [A management application client performs a discover process to discover all the SAS devices and expander devices in the SAS domain \(i.e., determining their device types, SAS addresses, and supported protocols\).](#) A SAS initiator device uses this information to determine SAS addresses to which it is able to establish connections. A self-configuring expander device uses this information to fill in its expander route table. ~~The discover process is the process of:~~

- ~~a) discovering all the SAS devices and expander devices in the SAS domain (i.e., determining their device types, SAS addresses, and supported protocols); and~~
- ~~b) configuring routing tables in the expander devices as needed.~~

A management application client performing the discover process shall perform a level-order (i.e., breadth-first) traversal of the SAS domain. The order of traversal shall be to discover:

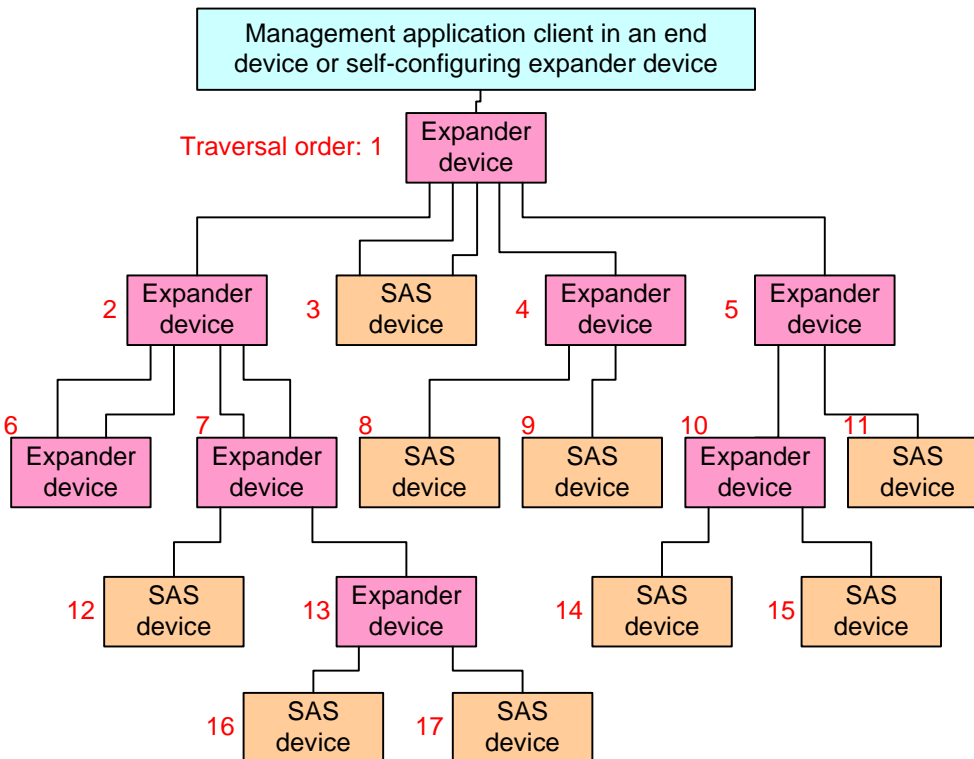
- 1) the device(s) to which the device containing the management application client is attached;
- 2) if an attached device is an expander device, every device attached to that expander device; and
- 3) for each expander device found, every device attached to that expander device.

This order is repeated until all expander devices have been traversed. [If the management application client discovers an externally configurable expander device that is not located beyond a self-configuring expander device with the CONFIGURES OTHERS bit set to one in the REPORT GENERAL response \(see 10.4.3.3\), it shall perform the configuration subprocess \(see 4.8\) to configure the expander route table before attempting to establish connections with devices attached two levels \(see 4.8.4\) beyond that externally configurable expander device.](#)

If an end device is directly attached to a self-configuring expander device with the CONFIGURES OTHERS bit set to one in the REPORT GENERAL response (see 10.4.3.3), then the management application client in that end device is not required to perform the configuration subprocess. If all the expander devices in the SAS domain are self-configuring expander devices, then management application clients in end devices are not required to perform the configuration subprocess.

If the management application client is inside a self-configuring expander device, then the process shall be repeated on each expander port.

Figure 12 shows an example of level-order traversal.



Note: Assume that the phy with the lowest phy identifier in each expander device is on the top right, and the remaining phys have increasing phy identifiers assigned in a counter-clockwise direction

Figure 12 — Level-order traversal example [changed to include “self-configuring”]

The management application client determines whether an expander device or SAS device is attached at each point in the traversal. For the first device (i.e., the device directly attached), this is determined from the DEVICE TYPE field in the IDENTIFY address frame (see 7.8.2) information received by the phy that the management application client is using. For other devices (i.e., devices not directly attached), this is determined from ATTACHED DEVICE TYPE field the SMP DISCOVER response (see 10.4.3.5).

If an expander device is attached, the management application client shall use the SMP REPORT GENERAL function (see 10.4.3.3) to determine how many phys are in the expander device and then use the SMP DISCOVER function (see 10.4.3.5) to determine what is attached to each expander phy (e.g., the device type, SAS address, and supported protocol(s)).

If the expander device’s EXTERNALLY CONFIGURABLE ROUTE TABLE bit is set to one in the SMP REPORT GENERAL response, the management application client shall configure its expander route table as described in 4.7.3 and 4.7.4.8.

If a self-configuring expander device’s CONFIGURING bit is set to one in the SMP REPORT GENERAL response, connection requests for destination ports two or more levels beyond the self-configuring expander device may result in responses such as OPEN_REJECT (NO DESTINATION) (see 7.2.5.12). If a management application client (e.g., in an end device) detects a self-configuring expander device with the

CONFIGURING bit set to one, the SAS port shall stop the I_T Nexus Loss timer for I_T nexuses involving such destination ports as long as the CONFIGURING bit is set to one (see 8.2.2.1).

If a SAS device is attached, the discover process is not required to obtain any more information about the SAS device. Additional discovery software may access that SAS device, however:

- a) if the SAS device supports an SMP target port, the management application client may use SMP functions (e.g., REPORT GENERAL and REPORT MANUFACTURER INFORMATION) to determine additional information about the SAS target device;
- b) if the SAS device supports an SSP target port, a SCSI application client may transmit SCSI commands (e.g., INQUIRY and REPORT LUNS) to determine additional information about the SCSI target device; and
- c) if the end device supports an STP target port, an ATA application client may transmit ATA commands (e.g., IDENTIFY DEVICE and IDENTIFY PACKET DEVICE) to determine additional information about the ATA device.

The result of the discover process is that the management application client has the necessary information (e.g., the device type, SAS address, and supported protocol(s)) to communicate with each SAS device and expander device in the SAS domain and each [externally](#) configurable expander device is configured with the necessary expander route entries to allow routing of connection requests through the SAS domain.

Annex L contains an example implementation of how a management application client may perform the discover process.

The discover process may be aborted prior to completion and restarted if there is an indication that it may be based on incorrect information (e.g., arrival of a Broadcast (Change)).

4.7.5 Discover process in a self-~~Self~~-configuring expander device [\[moved from below\]](#)

The management application client of a self-configuring expander device shall configure:

- a) the expander routing table in that expander device; and
- b) the expander routing table in each externally configurable expander device in the SAS domain that is not located behind another self-configuring expander device.

When a self-configuring expander device receives a Broadcast (Change) the management application client shall start the discover process using the expander port that received the Broadcast (Change). If a change to the expander route table is identified then the management device server shall set its CONFIGURING bit to one in the [SMP](#) REPORT GENERAL response (see 10.4.3.3).

The management application client shall set the CONFIGURING bit to zero when the discover process is complete. When the CONFIGURING bit changes from one to zero:

- a) a zoning expander device shall originate a Broadcast (Change) on each expander port that has access to the expander port through which the discover process was performed based on the zone permission table; and
- b) a non-zoning expander device shall originate a Broadcast (Change) on each expander port other than the one through which the discover process was performed.

When a Broadcast (Change) has been received by a self-configuring expander device, all previously valid SAS addresses shall continue to be routable until they are determined to be no longer valid. When a change has been identified, all unaffected SAS addresses shall continue to be routable.

[4.8 Configuration subprocess](#)

[4.8.1 Configuration subprocess overview](#)

[As part of the discover process \(see 4.7\), when it discovers externally configurable expander device, the management application client performs the configuration subprocess to configure the expander routing table in that externally configurable expander device with SAS addresses discovered two levels beyond each table routing phy in that externally configurable expander device. A single discover process performs the configuration subprocess at least once per externally configurable expander device.](#)

[The routing table in an expander device needs to be configured before connections are able to be established with devices attached two levels beyond that expander device.](#)

4.7.24.8.2 Allowed topologies

The management application client shall allow the following attachments between expander phys:

- a) ~~edge expander phy with the subtractive routing attribute attached to an edge expander phy with the subtractive routing attribute;~~
- b) ~~edge expander phy with the subtractive routing attribute attached to an edge expander phy with the table routing attribute;~~
- c) ~~edge expander phy with the subtractive routing attribute attached to a fanout expander phy with the table routing attribute.~~

If the management application client detects ~~any other combination of expander phy attachment (e.g., expander phy with table routing attached to expander phy with table routing, or an expander phy with direct routing attached to an expander phy)~~an externally configurable expander device phy with the table routing attribute attached to an expander device phy with either the direct routing attribute or the table routing attribute, it shall report an error in a vendor-specific manner.

If the management application client detects an overflow of ~~the edge~~an expander route index, it shall report an error in a vendor-specific manner.

If ~~the discover process~~ route table optimization (see 4.7.3) is disabled and the management application client detects an expander route entry that references the SAS address of the expander device itself (i.e., self-reference), it shall disable the expander route entry by setting the DISABLE EXPANDER ROUTE ENTRY bit to one in the SMP CONFIGURE ROUTE INFORMATION function request (see 10.4.3.13). The management application client shall disable each expander route entry in the expander route table by setting the DISABLE EXPANDER ROUTE ENTRY bit to one in the SMP CONFIGURE ROUTE INFORMATION function request (see 10.4.3.13) for each expander phy that has its attached device type set to 000b (i.e., no device attached).

If the management application client detects a port that the [discover configuration](#) process has not already configured with a SAS address it has already found attached to another expander device, it should use the SMP PHY CONTROL function (see 10.4.3.14) to disable all the expander phys attached to that SAS address except for phys in the expander device with the lowest SAS address. Figure 13 shows some illegal topologies.

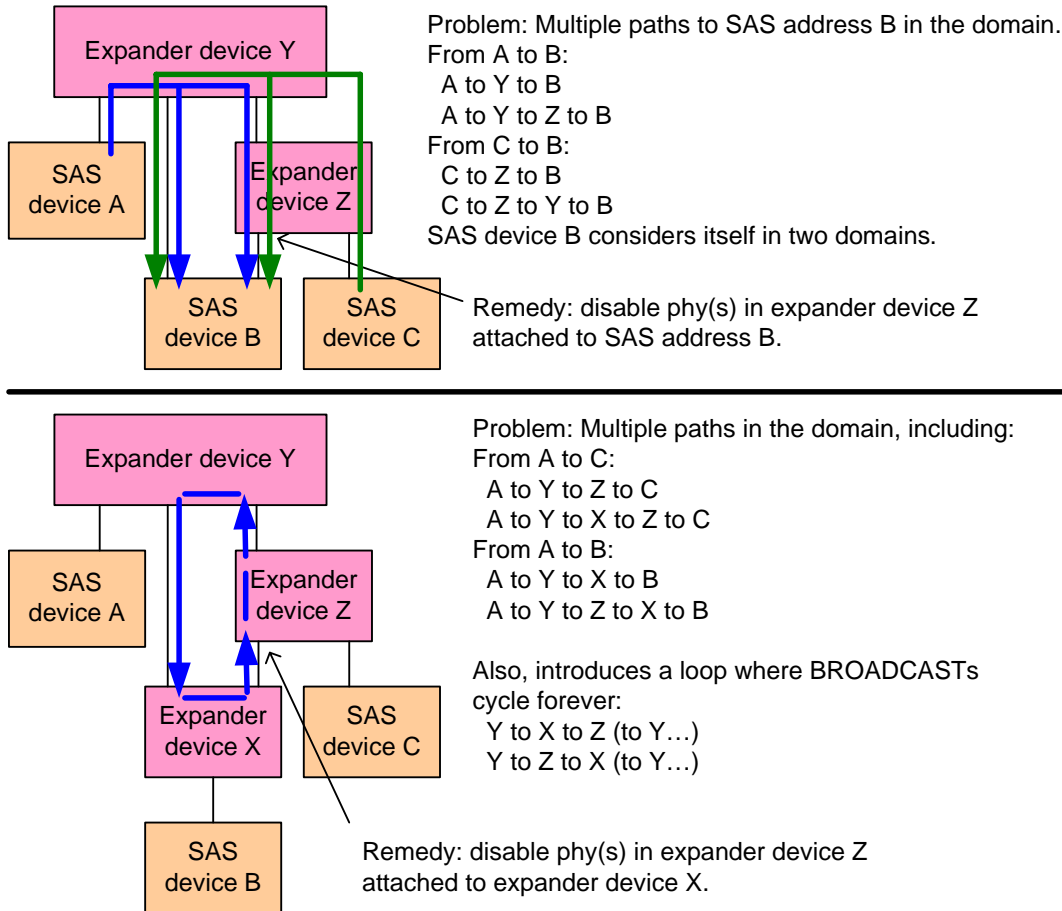


Figure 13 — Examples of invalid topologies

[Annex L](#) contains an example implementation of how a management application client may perform the [discover and configuration](#) process.

4.7.3.4.8.3 Discover process-Route table optimization

The management application client shall support a [discover process route table](#) optimization that reduces the number of entries required in an expander route table [in an externally configurable expander device](#). The method used to enable and disable [the discover process route table](#) optimization is vendor specific.

If [the discover process route table](#) optimization is enabled, then the management application client shall exclude discovered SAS addresses from the expander route table when any of the following conditions are met:

- a) in the discovered phy SMP DISCOVER response (see 10.4.3.5):
 - A) the FUNCTION RESULT field is set to a non-zero value (i.e., not SMP FUNCTION ACCEPTED);
- b) in the discovered phy SMP DISCOVER response:
 - A) the FUNCTION RESULT field is set to zero (i.e., SMP FUNCTION ACCEPTED);
 - B) the ROUTING ATTRIBUTE field is set to 1h (i.e., subtractive) or 2h (i.e., table); and
 - C) the ATTACHED DEVICE TYPE field is set to zero (i.e., no device attached);
- c) in the discovered phy SMP DISCOVER response:
 - A) the FUNCTION RESULT field is set to zero (i.e., SMP FUNCTION ACCEPTED);
 - B) the ROUTING ATTRIBUTE field is set to 1h (i.e., subtractive) or 2h (i.e., table);

- C) the ATTACHED DEVICE TYPE field is set to a non-zero value (e.g., end device, ~~edge expander device,~~ or ~~fanout~~ expander device); and
 - D) the ATTACHED SAS ADDRESS field contains the SAS address of the expander device being configured (i.e., a self-referencing address);
- d) in the discovered phy SMP DISCOVER response:
- A) the FUNCTION RESULT field is set to zero (i.e., SMP FUNCTION ACCEPTED);
 - B) the ROUTING ATTRIBUTE field is set to 1h (i.e., subtractive) or 2h (i.e., table);
 - C) the ATTACHED DEVICE TYPE field is set to a non-zero value (e.g., end device, ~~edge expander device,~~ or ~~fanout~~ expander device); and
 - D) the ATTACHED SAS ADDRESS field contains the SAS address of a device directly attached to the expander device being configured;
- or
- e) in the discovered phy SMP DISCOVER response:
- A) the FUNCTION RESULT field is set to zero (i.e., SMP FUNCTION ACCEPTED);
 - B) the ROUTING ATTRIBUTE field is set to 1h (i.e., subtractive) or 2h (i.e., table);
 - C) the ATTACHED DEVICE TYPE field is set to a non-zero value (e.g., end device, ~~edge expander device,~~ or ~~fanout~~ expander device); and
 - D) the ATTACHED SAS ADDRESS field contains a SAS address that already exists in the expander route table.

If the discovered SAS address being included in the expander route table is for a device that is currently not attached (i.e., the SMP DISCOVER response contains an ATTACHED DEVICE TYPE field set to zero (i.e., no device attached) and the ROUTE ATTRIBUTE field set to 0h (i.e., direct)), then the entry shall be inserted with the ROUTED SAS ADDRESS field set to 00000000 00000000h and the DISABLE EXPANDER ROUTE ENTRY bit set to one (see 10.4.3.13).

If ~~the discover process route table~~ optimization is disabled, then all SAS addresses shall be qualified for insertion in the expander route table.

If the management application client supports ~~the discover process route table~~ optimization, then the management application client should provide a vendor-specific method for initiating a check of the resulting expander route tables. The check should be performed under the following situations:

- a) when an OPEN_REJECT (NO DESTINATION) is received for a connection request to a SAS address that is expected to be accepted;
- b) when a discover process has been completed;
- c) when another SMP initiator port is discovered in the SAS domain; or
- d) when a self-configuring expander device is discovered in the SAS domain.

If the management application client detects an inconsistency in the expander route tables when ~~the discover process route table~~ optimization is enabled (e.g., detects entries that appear to have been filled in by a discover process with route table optimization disabled), then the management application client shall report an error in a vendor-specific manner and shall disable ~~the discover process route table~~ optimization. The management application client should then re-initiate a discover process with ~~the discover process route table~~ optimization disabled.

4.7.4.8.4 Expander route index order

The expander route table shall be configured for each expander phy that has a table routing attribute.

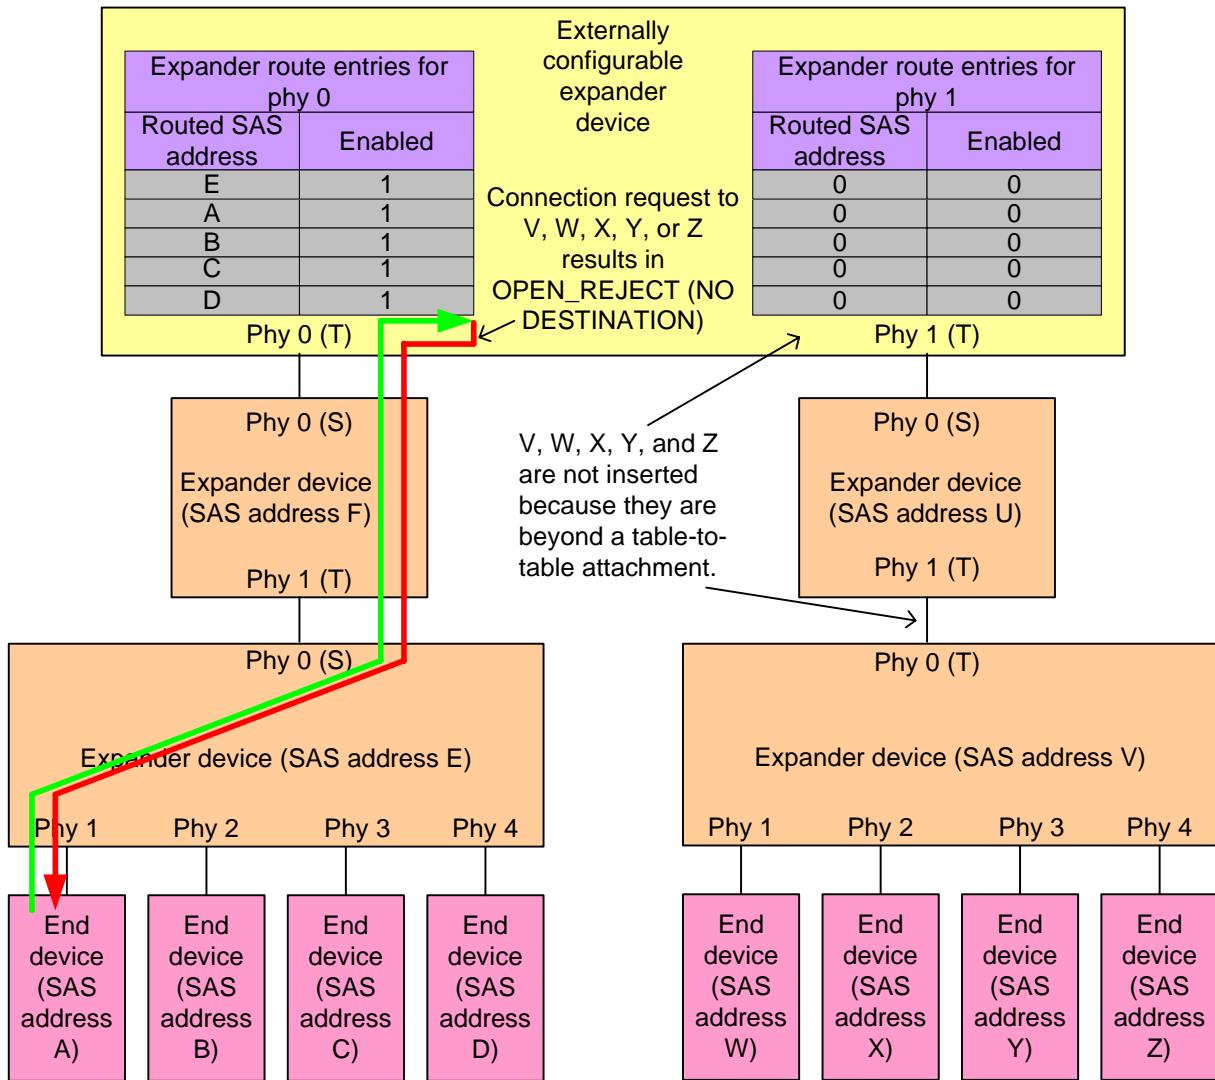
If the phy is not attached to an **edge** expander device, every expander route entry for that phy shall be disabled (i.e., the ROUTED SAS ADDRESS field shall be set to 00000000 00000000h and the DISABLE EXPANDER ROUTE ENTRY bit shall be set to one).

If the phy is attached to an **edge** expander device, the expander route table shall be configured for that phy as follows. For purposes of configuring the expander route table for that phy, the **edge** expander devices attached to the expander phy are assigned levels:

- 1) the expander device in which the expander route table is being configured is level 0;
- 2) the attached **edge** expander device is considered level 1;

- 3) devices attached to the level 1 **edge** expander device, except for the level 0 expander device, are considered level 2;
- 4) devices attached to level 2 **edge** expander devices, except for level 1 **edge** expander devices, are considered level 3; and
- 5) for each n greater than 3, devices attached to level n-1 **edge** expander devices, except for level n-2 **edge** expander devices, are considered level n.

SAS addresses of device attached beyond expander phys that are attached table-to-table shall not be included in the route table. Figure 14 shows an example of a route table that does not include SAS addresses beyond a table-to-table attachment.



Routing method
 S = subtractive
 T = table

Figure 14 — Externally configurable expander device and table-to-table attachment

NOTE 1 - This provides compatibility with management application clients compliant with previous versions of this standard. End devices in SAS domains containing externally configurable expander devices and table-to-table attachments may not be able to establish connections to each other. End devices in SAS domains containing only self-configuring expander devices compliant with this standard (i.e., supporting table-to-table attachments) are able to establish connections to any other end device.

Editor's Note 3: Verify that V is not added by any implementations. If the T-T check is only done while walking the phys of a device behind the T-T link (i.e., expander device V, it would not be.

The expander route table for each expander phy shall be configured starting from expander route index 0 by level (i.e., all level 1 entries first, then all level 2 entries, then all level 3 entries, etc.) up to the value of the EXPANDER ROUTE INDEXES field reported by the SMP REPORT GENERAL function (see 10.4.3.3).

Assuming the level 1 **edge** expander device has expander phys attached to N phys with qualified SAS addresses (see 4.7.3), the first N entries shall be used for those SAS addresses in expander phy order (i.e., the addresses attached to lower expander phy numbers first).

For each of the level 2 devices that:

- a) is an **edge** expander device attached to M phys with qualified SAS addresses; and
- b) is attached to an expander phy in the level 1 **edge** expander device with the table routing attribute,

the next M entries shall be used for the level 2 **edge** expander device's qualified SAS addresses in expander phy order (i.e., lower phy numbers first).

This process shall repeat for all levels of **edge** expander devices ~~in the edge expander device set~~.

After the expander route table has been configured with entries for all levels of edge expander devices, all remaining expander route **table** entries, if any, shall be disabled (i.e., the ROUTED SAS ADDRESS field shall be set to 00000000 00000000h and the DISABLE EXPANDER ROUTE ENTRY bit shall be set to one). The management application client is not required to disable entries if the topology of expander devices has not changed.

Figure 15 shows a portion of an edge expander device set SAS domain, where phy A in edge expander device R is being configured.

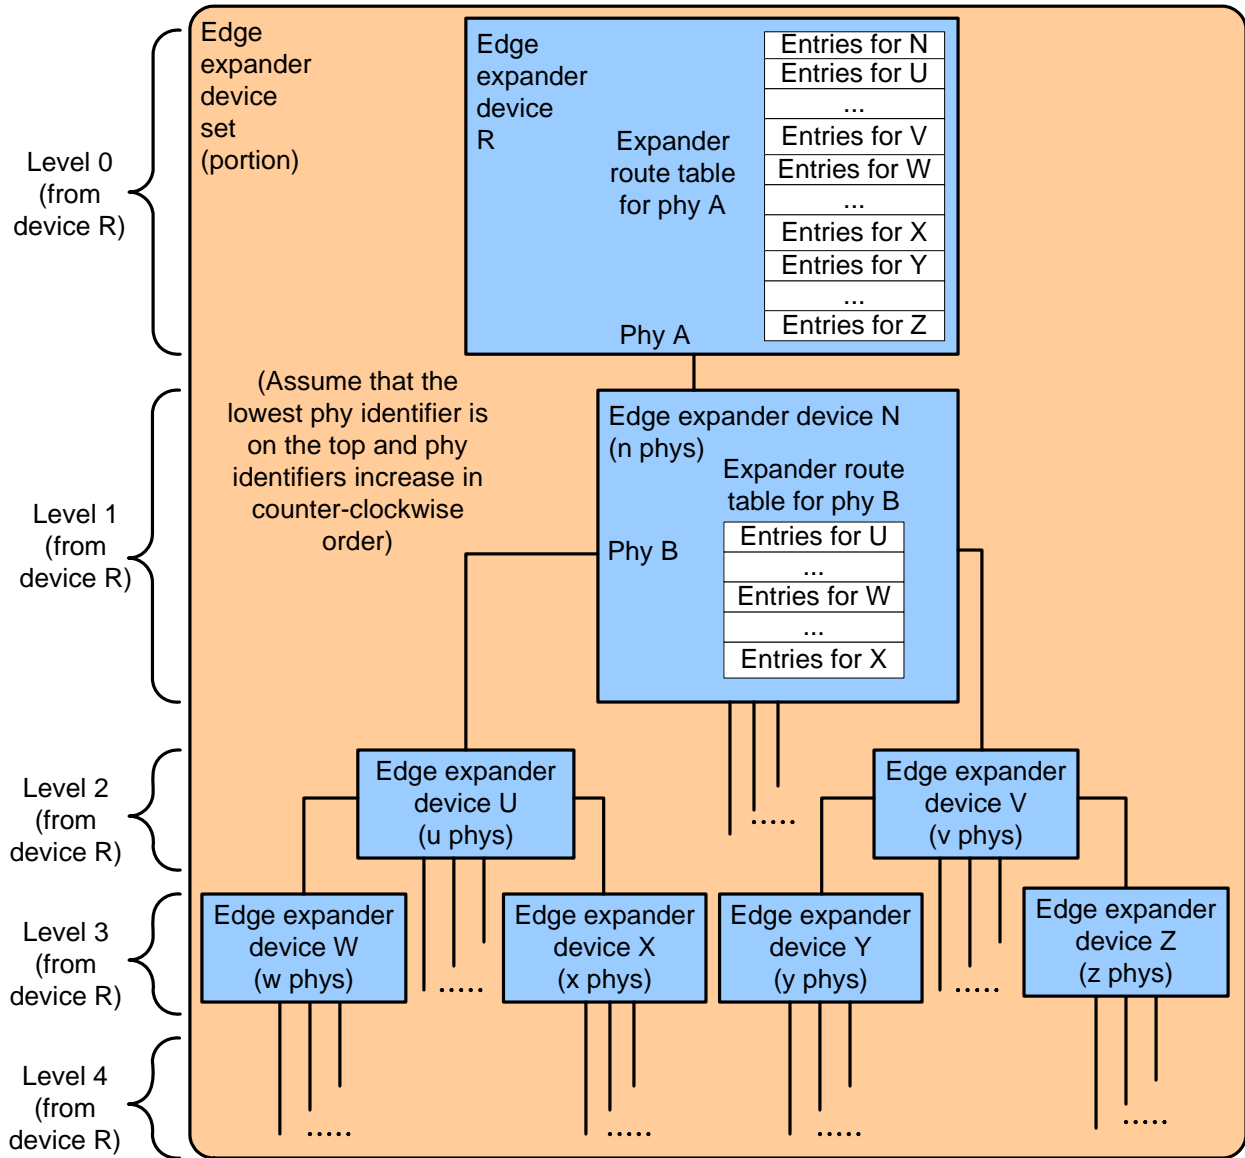


Figure 15 — Expander route index levels example [TODO delete "edge"]

Figure 16 shows a ~~fanout expander device and an edge expander device set~~, SAS domain where phy A in the ~~fanout~~ expander device R is being configured.

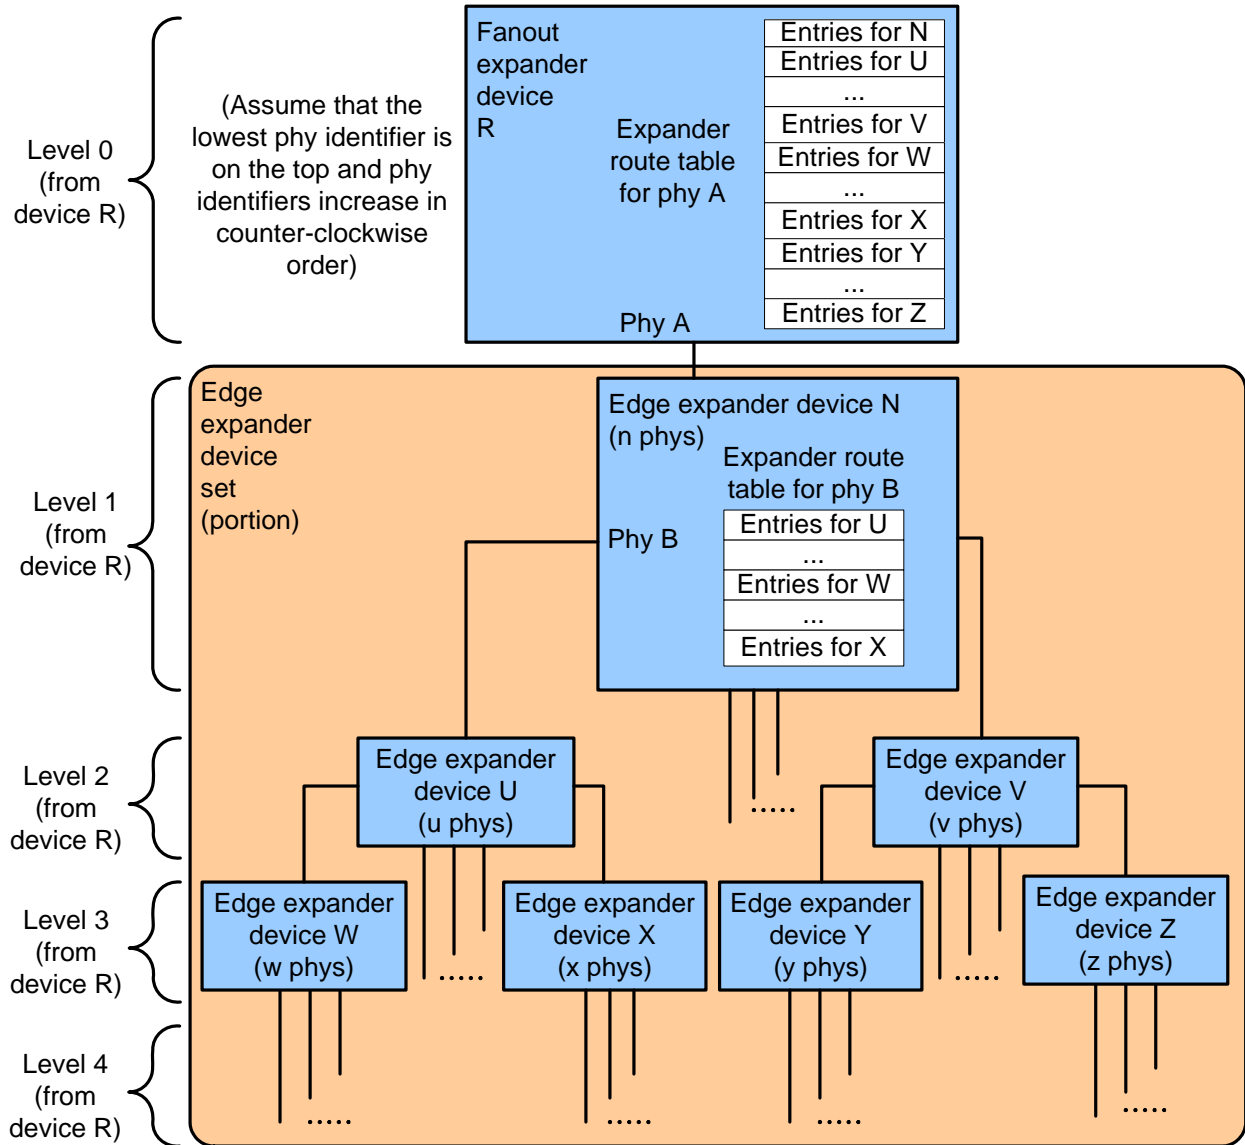


Figure 16 — Expander route index levels example with ~~fanout~~ expander device [TODO update]

Table 2 shows how the expander route table is configured for **edge externally configurable** expander device R phy A in figure 15 and the **fanout** expander device R phy A in figure 16.

Table 2 — Expander route table levels for **edge externally configurable expander device R-~~or fanout~~ expander device R**

Expander route index	Expander route entry contents
Level 1 (from device R) entries	
0 .. ($\leq n$ entries)	Qualified SAS addresses attached to edge expander device N
Level 2 (from device R) entries	
($\leq u$ entries)	Qualified SAS addresses attached to edge expander device U
...	...additional qualified SAS addresses for expander devices at level 2...
($\leq v$ entries)	Qualified SAS addresses attached to edge expander device V
Level 3 (from device R) entries	
($\leq w$ entries)	Qualified SAS addresses attached to edge expander device W
...	...additional qualified SAS addresses for expander devices at level 3...
($\leq x$ entries)	Qualified SAS addresses attached to edge expander device X
($\leq y$ entries)	Qualified SAS addresses attached to edge expander device Y
...	...additional qualified SAS addresses for expander devices at level 3...
($\leq z$ entries)	Qualified SAS addresses attached to edge expander device Z
Entries for additional levels	
...	...
Disabled entries	
...	...

Table 3 shows how the expander route table is configured for **edge externally configurable** expander device N phy B in figure 15 and figure 16.

Table 3 — Expander route table levels for **edge externally configurable expander device N**

Expander route index	Expander route entry contents
Level 1 (from device N) entries	
($\leq u$ entries)	Qualified SAS addresses attached to edge expander device U
Level 2 (from device N) entries	
($\leq w$ entries)	Qualified SAS addresses attached to edge expander device W
...	...additional qualified SAS addresses for expander devices at level 2...
($\leq x$ entries)	Qualified SAS addresses attached to edge expander device X
Entries for additional levels	
...	...
Disabled entries	
...	...

Figure 17 shows an example topology ~~with a fanout expander device.~~

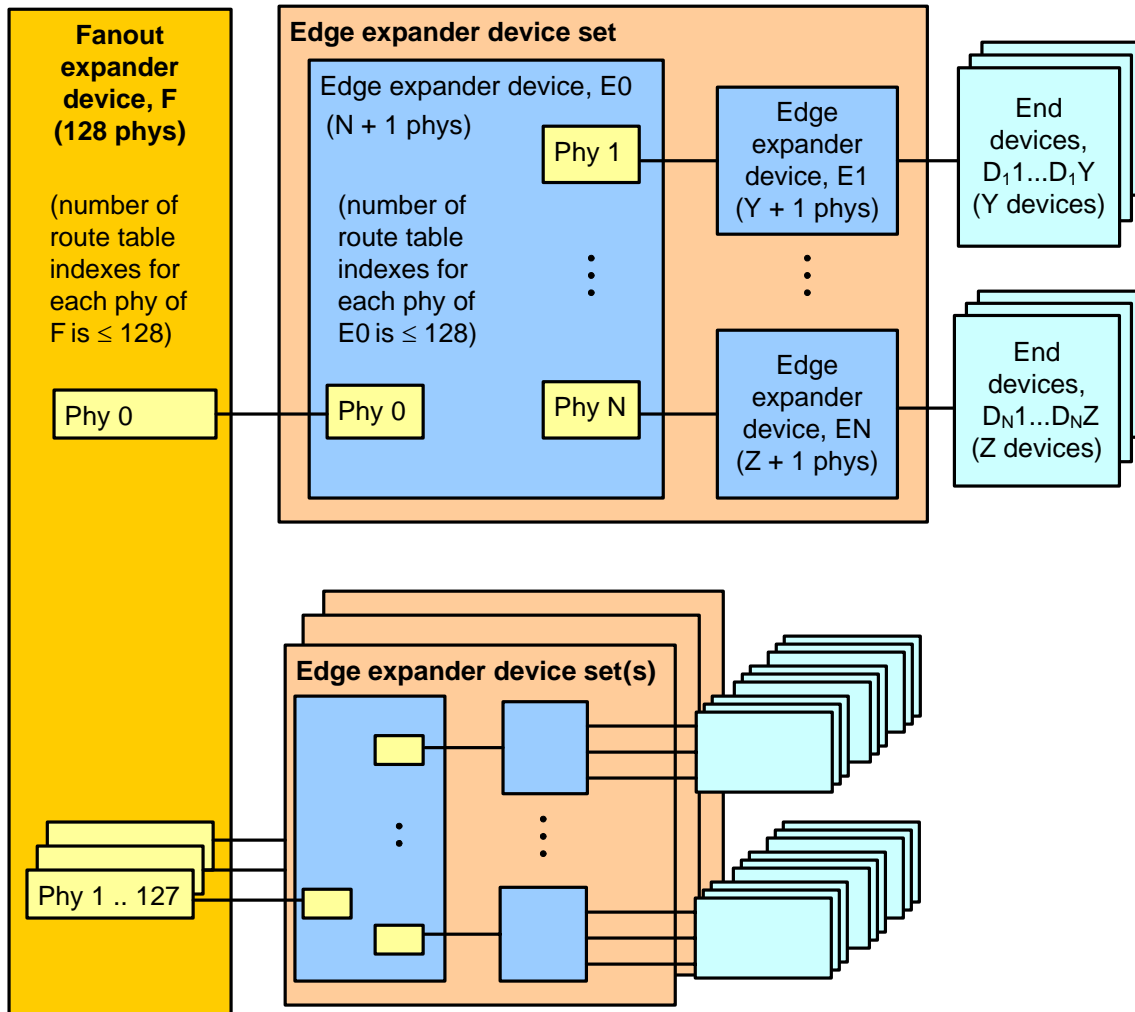


Figure 17 — Expander route index order example

Table 4 shows the expander route index order for ~~the edge~~ expander E0 phy 1 in figure 17.

Table 4 — Expander route entries for ~~edge~~ expander E0 phy 1

Expander route index	Expander route entry contents
Level 1 entries	
0	SAS address (e.g., D ₁ 1) of the device attached to phy 1 of edge expander device E1
1	SAS address (e.g., D ₁ 2) of the device attached to phy 2 of edge expander device E1
...	...
Y - 1	SAS address (e.g., D ₁ Y) of the device attached to phy Y of edge expander device E1
Level 2 and beyond: no entries	
Disabled entries	

Table 5 shows the expander route index order for ~~the fanout~~ expander F phy 0 in figure 17.

Table 5 — Expander route entries for ~~fanout~~ expander device F phy 0

Expander route index	Expander route entry contents
Level 1 entries	
0	SAS address (e.g., E1) of the device attached to phy 1 of edge expander device E0
...	...additional qualified SAS addresses for edge expander device E0...
N - 1	SAS address (e.g., EN) of the device attached to phy N of edge expander device E0
Level 2 entries	
N	SAS address (e.g., D ₁ 1) of the device attached to phy 1 of edge expander device E1
...	...additional qualified SAS addresses for edge expander device E1...
	SAS address (e.g., D ₁ Y) of the device attached to phy Y of edge expander device E1
...	...additional qualified SAS addresses for edge expander devices E2 through EN-1...
	SAS address (e.g., D _N 1) of the device attached to phy 1 of edge expander device EN
...	...additional qualified SAS addresses for edge expander device EN...
	SAS address (e.g., D _N Z) of the device attached to phy Z of edge expander device EN
Level 3 and beyond: no entries since all devices attached to E1 through EN, except for E0, are end devices	
Disabled entries	

Suggested changes to chapter 5 (Physical layer)

5.2.3 Connectors

5.2.3.3 SAS external connectors

5.2.3.3.1 SAS 4x connectors

5.2.3.3.1.1 SAS 4x cable plug connector

...

Table 6 defines the icons that should be placed on or near SAS 4x cable plug connectors.

Table 6 — SAS 4x cable plug connector icons

Use	Icon
End of a SAS external cable that attaches to an end device or an enclosure out port or an enclosure universal port	Diamond
End of a SAS external cable that attaches to an end device or an enclosure in port or an enclosure universal port	Circle

5.2.3.3.1.2 SAS 4x receptacle connector

...

Table 7 defines the icons that should be placed near SAS 4x receptacle connectors.

Table 7 — SAS 4x receptacle connector icons

Use	Icons
Enclosure out port (see 4.6.2)	Diamond
End device <u>or enclosure universal port</u>	Diamond and circle
Enclosure in port (see 4.6.2)	Circle

5.2.3.3.2 Mini SAS 4x connectors

5.2.3.3.2.1 Mini SAS 4x cable plug connector

...

Table 8 defines the icons that shall be placed on or near Mini SAS 4x cable plug connectors and the key slot positions (see SFF-8088) that shall be used by Mini SAS 4x cable plug connectors.

Table 8 — Mini SAS 4x cable plug connector icons and key slot positions

Use	Icon	Key slot positions	Reference
End of a SAS external cable that attaches to an end device <u>or an enclosure out port</u> , <u>or an enclosure universal port</u>	Diamond	2, 4	Figure 18
End of a SAS external cable that attaches to an end device <u>or an enclosure in port</u> , <u>or an enclosure universal port</u>	Circle	4, 6	Figure 19

Figure 18 shows the key slots on the Mini SAS 4x cable plug connector that attaches to either an end device (see figure 20 in 5.2.3.3.2.2) or an enclosure out port (see figure 21 in 5.2.3.3.2.2) or an enclosure universal port.




Figure 18 — Mini SAS 4x cable plug connector that attaches to an enclosure out port

Figure 19 shows the key slots on the Mini SAS 4x cable plug connector that attaches to either an end device (see figure 20 in 5.2.3.3.2.2), ~~or~~ an enclosure in port (see figure 22 in 5.2.3.3.2.2), or an enclosure universal port.




Figure 19 — Mini SAS 4x cable plug connector that attaches to an enclosure in port

5.2.3.3.2.2 Mini SAS 4x receptacle connector

...

Table 9 defines the icons that shall be placed on or near Mini SAS 4x receptacle connectors and the key positions (see SFF-8088) that shall be used by Mini SAS 4x receptacle connectors.

Table 9 — Mini SAS 4x receptacle connector icons and key positions

Use	Icons	Key position	Reference
End device <u>or enclosure universal port</u>	Diamond and circle	4	Figure 20
Enclosure out port (see 4.6.2)	Diamond	2	Figure 21
Enclosure in port (see 4.6.2)	Circle	6	Figure 22

Figure 20 shows the key on a Mini SAS 4x receptacle connector used by end devices. This connector may be attached to the Mini SAS 4x cable plug shown in figure 18 or the Mini SAS 4x cable plug shown in figure 19 (see 5.2.3.3.2.2).




Figure 20 — Mini SAS 4x receptacle connector - end device or enclosure universal port

Figure 21 shows the key on a Mini SAS 4x receptacle connector used by an enclosure out port. This connector may be attached to the Mini SAS 4x cable plug shown in figure 18 (see 5.2.3.3.2.2).




Figure 21 — Mini SAS 4x receptacle connector - enclosure out port

Figure 22 shows the key on a Mini SAS 4x receptacle connector used by an enclosure in port. This connector may be attached to the Mini SAS 4x cable plug shown in figure 19 (see 5.2.3.3.2.2).




Figure 22 — Mini SAS 4x receptacle connector - enclosure in port

5.2.4 Cable assemblies

5.2.4.2 SAS external cable assemblies

5.2.4.2.1 SAS external cable assemblies overview

There are several types of SAS external cable assemblies:

- SAS 4x cable plug connector (see 5.2.3.3.1.1) at each end (see 5.2.4.2.2);
- Mini SAS 4x cable plug connector (see 5.2.3.3.2.1) at each end (see 5.2.4.2.3); and
- SAS 4x cable plug connector at one end and Mini SAS 4x cable plug connector at the other end (see 5.2.4.2.4).

SAS external cable assemblies do not include power or the READY LED signal.

Although the connector always supports four physical links, a SAS external cable assembly may support one, two, three, or four physical links. SAS external cable assemblies should be labeled to indicate how many physical links are included (e.g., 1X, 2X, 3X, and 4X on each connector's housing)

The Tx signals on one end shall be connected to the corresponding Rx signals of the other end (e.g., Tx 0+ of one connector shall be connected to Rx 0+ of the other connector).

Signal returns shall not be connected to CHASSIS GROUND in the cable assembly.

In addition to the SAS icon (see Annex M), additional icons are defined for external connectors to guide users into making compatible attachments (i.e., not attaching expander device table routing phys to expander device table routing phys [in externally configurable expander devices](#), which is not allowed by this standard). Connectors that have one or more matching icons are intended to be attached. Connectors that do not have a matching icon should not be attached together.

One end of the SAS external cable assembly shall support being attached to an end device, ~~or~~ an enclosure out port, [or an enclosure universal port](#). The other end of the SAS external cable assembly shall support being attached to an end device, ~~or~~ an enclosure in port, [or an enclosure universal port](#). If a SAS 4x cable plug connector is used, it should include icons as defined in 5.2.3.3.1.1. If a Mini SAS 4x cable plug connector is used, it shall include icons and key slots as defined in 5.2.3.3.2.1.

Suggested changes to chapter 7 (Link layer)

7.8.2 IDENTIFY address frame

...

The DEVICE TYPE field specifies the type of device containing the phy, and is defined in table 10.

Table 10 — DEVICE TYPE field

Code	Description
001b	End device
010b	Edge -expander device
011b	Fanout e Expander device compliant with a previous version of this standard
All others	Reserved

7.9.2 SAS initiator device rules

After a link reset sequence, or after receiving a BROADCAST (CHANGE), a management application client behind an SMP initiator port should perform a discover process (see 4.7).

When a discover process is performed after a link reset sequence, the management application client discovers all the devices in the SAS domain. When a discover process is performed after a BROADCAST (CHANGE), the management application client determines which devices have been added to or removed from the SAS domain.

The discover information may be used to select connection rates for connection requests (see 7.8.3).

7.9.3 ~~Fanout e~~Expander device rules

After completing the identification sequence on a phy and completing internal initialization, the ECM within an ~~fanout~~ expander device shall be capable of routing connection requests through that phy. The expander device may return OPEN_REJECT (NO DESTINATION) until it is ready to process connection requests.

After a link reset sequence, or after receiving a BROADCAST (CHANGE), the management application client behind an SMP initiator port in a ~~fanout self-configuring~~ expander device ~~that does not have a configurable expander route table~~ shall follow the SAS initiator device rules (see 7.9.2) to perform a discover process ([see 4.7](#)).

The ECM of a ~~fanout externally configurable~~ expander device ~~that has a configurable expander route table~~ is dependent on the completion of the discover process (see 4.7) for routing connection requests using the table routing method.

~~7.9.4 Edge expander device rules~~

~~After completing the identification sequence on a phy and completing internal initialization, the ECM within an edge expander device shall be capable of routing connection requests through that phy. The expander device may return OPEN_REJECT (NO DESTINATION) until it is ready to process connection requests.~~

~~The ECM of an edge expander device that has a configurable expander route table is dependent on the completion of the discover process (see 4.7) for routing connection requests using the table routing method.~~

7.10 Power management

...

ATA idle and standby power modes, implemented with the IDLE, IDLE IMMEDIATE, STANDBY, STANDBY IMMEDIATE, and CHECK POWER MODE commands (see ~~ATA/ATAPI-7 V1~~ATA8-ACS), may be supported by STP initiator ports. The ATA sleep power mode, implemented with the SLEEP command, shall not be used.

7.11 SAS domain changes

After power on or receiving Broadcast (Change), the management application client is ~~in~~behind each SAS initiator port should scan the SAS domain using the discover process (see 4.7) to search for SAS initiator devices, SAS target devices, and expander devices.

The expander device shall originate Broadcast (Change) from at least one phy in each expander port other than the expander port that is the cause for originating Broadcast (Change).

Expander devices shall originate Broadcast (Change) for the following reasons:

- a) after an expander phy's SP state machine transitions from the SP15:SAS_PHY_Ready or SP22:SATA_PHY_Ready state to the SP0:OOB_COMINIT state (see 6.8);

NOTE 2 - This occurs when the expander phy is reset or disabled with the SMP PHY CONTROL function DISABLE, LINK RESET, HARD RESET, or TRANSMIT SATA PORT SELECTION SIGNAL phy operations (see 10.4.3.14) as well as when dword synchronization is unexpectedly lost;

- b) after a virtual phy has been disabled with the SMP PHY CONTROL function DISABLE phy operation or started processing a reset requested by the LINK RESET or HARD RESET phy operations (see 10.4.3.14);
- c) after an expander phy's SP state machine reaches the SATA spinup hold state (see 6.8.7 and 6.10);
- d) after an expander phy's SP state machine receives a COMWAKE Detected message in states SP0:OOB_COMINIT, SP1:OOB_AwaitCOMX, SP3:OOB_AwaitCOMINIT_Sent, or SP4:OOB_COMSAS if ~~the value of~~ the ATTACHED SATA PORT SELECTOR bit is set to zero in the DISCOVER response prior to receiving the COMWAKE detected message (see 6.8.3 and table 207 in 10.4.3.5);
- e) after an expander phy's SP state machine transitions from the SP1:OOB_AwaitCOMX state to the SP0:OOB_COMINIT state if ~~the value of~~ the ATTACHED SATA PORT SELECTOR bit was set to one in the DISCOVER response upon entry to SP1:OOB_AwaitCOMX, and if no COMWAKE detected message was received while in SP1:OOB_AwaitCOMX before the transition to SP0:OOB_COMINIT (see 6.8.3.3.2);
- f) after the link reset sequence completes (see 7.9);
- g) after a virtual phy has been enabled or completed processing a reset requested by the SMP PHY CONTROL function LINK RESET or HARD RESET phy operations (see 10.4.3.14);
- h) after a self-configuring expander device has changed its CONFIGURING bit from one to zero in the SMP REPORT GENERAL response (see 10.4.3.3) as described in 4.7.5;
- i) after an STP/SATA bridge receives an initial Register - Device to host FIS (see 9.3.1).

Expander devices shall forward Broadcast (Change) for the following reasons:

- a) after an expander phy receives BROADCAST (CHANGE).

For a virtual phy, if there is any time after a reset is originated during which connection requests to the attached SAS address result in connection responses of OPEN_REJECT (NO DESTINATION), the expander device shall originate the Broadcast (Change) twice, once at the start of the reset (i.e., when the SAS address becomes unavailable) and once at its completion (i.e., when the SAS address becomes available). If there is no such time window, the expander device shall originate the Broadcast (Change) once.

SAS initiator ports may originate Broadcast (Change) to force other SAS initiator ports and expander ports to re-run the discover process, but should not be sent by SAS target ports.

A SAS initiator port that receives Broadcast (Change) shall follow the SAS initiator device rules (see 7.9.2) to discover [and configure](#) the topology.

An [fanout](#) expander device that receives Broadcast (Change) shall follow the [fanout expander](#) device rules (see 7.9.3) to discover [and configure](#) the topology.

~~An edge expander device that receives Broadcast (Change) shall follow the edge device rules (see 7.9.4).~~

See 10.4.3.3 for details on counting Broadcast (Change) origination in an expander device.

7.12 Connections

7.12.4 Arbitration and resource management in an expander device

7.12.4.3 ~~Edge~~ Expander devices

When the ECM in an [edge](#) expander device receives a connection request:

- 1) if the destination SAS address is that of the expander device itself, the ECM shall arbitrate for access to its SMP target port and forward the connection request;
- 2) if the destination SAS address matches the SAS address of a device to which one of the expander phys is attached, the ECM shall arbitrate for access to one of the matching phys and forward the connection request;
- 3) if the destination SAS address matches an enabled SAS address in the expander route table for one or more expander phys which has a table routing attribute (see 4.6.7.1) and is attached to an [edge](#) expander device, the ECM shall arbitrate for access to one of the matching phys and forward the connection request; and
- 4) if at least one phy has the subtractive routing attribute and is attached to an expander device (i.e., ~~it is attached to an edge expander device or a fanout expander device and~~ the phy is using the subtractive routing method), and the request did not come from that expander device, the ECM shall arbitrate for access to one of the subtractive routing phys and forward the connection request.

If it does not find a match and no phy using the subtractive routing method exists, the ECM shall reply with Arb Reject (No Destination).

If the destination phy is in the same expander port as the source phy and is using the subtractive routing method, the ECM shall reply with Arb Reject (No Destination).

If the destination phy is in the same expander port as the source phy and is using the direct routing method or the table routing method, the ECM shall reply with either Arb Reject (No Destination) or Arb Reject (Bad Destination); it should reply with Arb Reject (No Destination).

NOTE 3 - NOTE 32 - ECMs in [edge](#) expander devices compliant with previous versions of this standard were required to reply with Arb Reject (Bad Destination).

~~7.12.4.4 Fanout expander devices~~

~~When the ECM in a fanout expander device receives a connection request:~~

- ~~1) if the destination SAS address is that of the expander device itself, the ECM shall arbitrate for access to its SMP target port and forward the connection request;~~
- ~~2) if the destination SAS address matches the SAS address of a device to which one of the expander phys is attached, the ECM shall arbitrate for access to one of the matching phys and forward the connection request; and~~
- ~~3) if the destination SAS address matches an enabled SAS address in the expander route table for one or more expander phys which has a table routing attribute (see 4.6.7.1) and is attached to an edge expander device, the ECM shall arbitrate for access to one of the matching phys and forward the connection request.~~

~~If it does not find a match, the ECM shall reply with Arb Reject (No Destination).~~

~~If the destination phy is in the same expander port as the source phy, the ECM shall reply with either Arb Reject (No Destination) or Arb Reject (Bad Destination); it should reply with Arb Reject (No Destination).~~

~~NOTE 4 – NOTE 33 – ECMs in fanout expander devices compliant with previous versions of this standard were required to reply with Arb Reject (Bad Destination).~~

Editor's Note 4: delete all cross references to 7.12.4.4 (in 4.6.6.3, 7.2.5.11, 7.12.4.1)

Suggested changes to chapter 10 (Application layer)

10.4.3.3 REPORT GENERAL function

The REPORT GENERAL function returns general information about the SAS device (e.g., a SAS device contained in an expander device). This SMP function shall be implemented by all SMP target ports.

...

Table 11 defines the response format.

Table 11 — REPORT GENERAL response

Byte\Bit	7	6	5	4	3	2	1	0
0	SMP FRAME TYPE (41h)							
1	FUNCTION (00h)							
2	FUNCTION RESULT							
3	RESPONSE LENGTH (08h)							
4	(MSB)	EXPANDER CHANGE COUNT						(LSB)
5								
6	(MSB)	EXPANDER ROUTE INDEXES						(LSB)
7								
8	Reserved							
9	NUMBER OF PHYS							
10	TABLE TO TABLE SUPPORTED	Reserved			CONFIGURES OTHERS	CONFIGURING	EXTERNALLY CONFIGURABLE ROUTE TABLE	
11	Reserved							
12	ENCLOSURE LOGICAL IDENTIFIER							
19								
20	Reserved							
29								
30	(MSB)	STP BUS INACTIVITY TIME LIMIT						(LSB)
31								
32	(MSB)	STP MAXIMUM CONNECT TIME LIMIT						(LSB)
33								
34	(MSB)	STP SMP I_T NEXUS LOSS TIME						(LSB)
35								
36	(MSB)	CRC						(LSB)
39								

The SMP FRAME TYPE field shall be set to 41h.

The FUNCTION field shall be set to 00h.

The FUNCTION RESULT field is defined in 10.4.3.2.

The RESPONSE LENGTH field shall be set to 08h. For compatibility with previous versions of this standard, a RESPONSE LENGTH field set to 00h specifies that there are 6 dwords before the CRC field.

The EXPANDER CHANGE COUNT field counts the number of Broadcast (Change)s originated by an expander device (see 7.11). SMP target ports in expander devices shall support this field. SMP target ports in other device types (e.g., end devices) shall set the EXPANDER CHANGE COUNT field to 0000h. This field shall be set to 0000h at power on. If the expander device has originated Broadcast (Change) for any reason described in 7.11 since transmitting a REPORT GENERAL response, it shall increment this field at least once from the value in the previous REPORT GENERAL response. It shall not increment this field when forwarding a Broadcast (Change). This field shall wrap to zero after the maximum value (i.e., FFFFh) has been reached.

NOTE 5 - Application clients that use the EXPANDER CHANGE COUNT field should read it often enough to ensure that it does not increment a multiple of 65 536 times between reading the field.

The EXPANDER ROUTE INDEXES field contains the maximum number of route indexes per phy for the expander device (see 4.6.7.3). SMP target ports in [externally configurable](#) expander devices shall support this field. SMP target ports in other device types (e.g., end devices [and self-configuring expander devices](#)) shall set the EXPANDER ROUTE INDEXES field to zero. Not all phys in a ~~n-edge~~ [externally configurable](#) expander device are required to support the maximum number indicated by this field.

The NUMBER OF PHYS field contains the number of phys in the device, including any virtual phys and any vacant phys.

[A TABLE TO TABLE SUPPORTED bit set to one indicates the expander device is a self-configuring expander device that supports its table routing phys being attached to table routing phys in other expander devices. The TABLE TO TABLE SUPPORTED bit shall only be set to one if the EXTERNALLY CONFIGURABLE ROUTE TABLE bit is set to zero. A TABLE TO TABLE SUPPORTED bit set to zero indicates the expander device is not a self-configuring expander device that supports its table routing phys being attached to table routing phys in other expander devices.](#)

[A CONFIGURES OTHERS bit set to one indicates that the expander device is a self-configuring expander device that performs the configuration subprocess defined in 4.8. A CONFIGURES OTHERS bit set to zero indicates the expander device may or may not perform the configuration subprocess. Self-configuring expander devices compliant with this standard shall set the CONFIGURES OTHERS bit to one.](#)

[NOTE 6 - If the CONFIGURES OTHERS bit is set to zero, the expander device may configure all externally configurable expander devices in the SAS domain.](#)

A CONFIGURING bit set to one indicates that ~~a self-configuring expander device has not completed configuring its expander route table~~ [the SMP target port is in a self-configuring expander device, its management application client is currently performing the discover process \(see 4.7\), and it has identified at least one change to its expander routing table.](#) A CONFIGURING bit set to zero indicates that ~~configuration is complete and the expander device is ready for connection requests~~ [the SMP target port is not in a self-configuring expander device currently performing the discover process and changing its expander routing table.](#) Changes in this bit from one to zero result in a Broadcast (Change) being originated (see 7.11). SMP target ports in self-configuring expander devices shall support this bit. SMP target ports in [externally](#) configurable expander devices and in other device types shall set the CONFIGURING bit to zero.

~~The~~ [A EXTERNALLY CONFIGURABLE ROUTE TABLE bit set to one indicates whether the that the SMP target port is in an externally configurable expander device and has an expander route table that is required to be configured with the SMP CONFIGURE ROUTE INFORMATION function \(see 4.6.7.3\). An expander device with a configurable route table shall set the CONFIGURABLE ROUTE TABLE bit to one. An EXTERNALLY CONFIGURABLE ROUTE TABLE bit set to zero indicates that the SMP target port is not in an externally configurable expander device \(e.g., it is in an end device, in a self-configuring expander device, or in an expander device with no phys with table routing attributes\) or a device with any other device type shall set the CONFIGURABLE ROUTE TABLE bit to zero.](#)

...

10.4.3.5 DISCOVER function

...

The ATTACHED DEVICE TYPE field indicates the DEVICE TYPE value received during the link reset sequence and is defined in table 10.

Table 12 — ATTACHED DEVICE TYPE field

Code	Description
000b	No device attached
001b	End device
010b	Edge expander device
011b	Fanout e Expander device compliant with a previous version of this standard
All others	Reserved

...

Suggested changes to annexes

Annex L (informative) Discover process example implementation

L.1 Discover process example implementation overview

This annex includes a C program implementing the discover process. Table L.1 describes the source files.

L.2 Header file

The following is the C header file for the discover process.

...

[Editor's Note 5: Deferred to another proposal: changes to the .h file are necessary to remove edge/fanout and edge expander device set handling](#)

L.3 Source file

The following is the C source file for the discover process.

...

[Editor's Note 6: Deferred to another proposal: changes to the .cpp file are necessary to remove edge/fanout and edge expander device set handling](#)
