

To: T10 Technical Committee
 From: Luben Tuikov, Vitesse Semiconductor (ltuikov@vitesse.com)
 Date: 20 October 2006
 Subject: 06-460r2 SAM-4 Time bounds of task states

Revision history

Revision 0 (18 October 2006) First revision

Revision 1 (19 October 2006) Second revision

- 1) Figure: named timeout periods.
- 2) Figure: Updated task ending to a non-trivial example.
- 3) Text: Expressed and clarified relationship between timeouts.
- 4) Text: Updated Summary section.
- 5) Text: Various other small corrections throughout the document.

Revision 2 (20 October 2006) Third revision

Additions in blue with a change bar, removals in red, ~~strikethrough~~ with a change bar.

- 1) **Suggested changes** section additions: SAM-4 inclusion intent of this document.
- 2) **Implications** section editing: **Ta** and **Tb** definition and relationship.
- 3) **Summary** section additions: intent of this proposal.

Related documents

sam4r07 - SCSI Architecture Model - 4, revision 07

spc4r07a - SCSI Primary Commands -4, revision 07a

06-368r0 - SAM-4 QUERY TASK TMF Extended Response, revision 0 (A)

05-284r3 - SPC-4: Self Describing Command Timeouts, revision 3 (B)

Overview

The SCSI Architecture Model document gives detailed description of task states and task set management in section 8.

Recently, we've seen the need for SCSI Targets to be able to convey to Application clients, a command timeout value such that Application clients can better manage error recovery and avoid error recovery "thrashing". See Related documents (A) and (B).

The task states defined in the SCSI Architecture Model are TASK DORMANT, TASK ENABLED, TASK BLOCKED and TASK ENDED. Task set management is defined in terms of the four task states mentioned here, four task attributes and a priority attribute which we will not mention here as it is not relevant to the purpose of this proposal.

That is, SAM-4 compliant task servers implementing task management and a task manager as specified in SAM-4, should at least be able to represent the four task states. In other words, they should be able to map from an internal (implementation and/or proprietary) task state to one of these SAM-4 defined four states.

Error Recognition and Error Recovery

At the application client, error recognition can happen in one of the following ways:

- a) Service response,
- b) Task status code, or
- c) Application client timed out the task after certain timeout period had been allocated for its execution.

Case c) is particularly interesting as it is the topic of 05-284r3 - SPC-4: Self Describing Command Timeouts.

After the task has timed-out its allocated execution time, the application client would normally issue an ABORT TASK TMF. This isn't always optimal. Depending on the task state at the device server (target), an application client may abort a perfectly good task, for example in TASK ENABLED or TASK ENDED state. For this reason, it is encouraged that application clients first issue QUERY TASK TMF in order to figure out where the task is and what its state is and then decide whether to issue an ABORT TASK TMF or whether to reset their task execution timeout timer and wait again for a certain amount of retries (finite), thus extending the total execution timeout.

This method works fine for some of the cases but not for all. For example, if the task is in TASK DORMANT state after it timed out the first time, an application client may wish to abort the task, and notify the system administrator that

external intervention may be required to the logical unit. If, on the other hand, the task is in TASK ENABLED state after it timed out for the first time, it is advised that application clients reset the timer and let the logical unit complete the task.

This is where proposal 06-368r0 - SAM-4 QUERY TASK TMF Extended Response, extends the response of QUERY TASK TMF to optionally to the device server, return the task state of the task being queried, so that application clients can decide whether to abort the task (e.g. TASK DORMANT) or give it another chance (TASK ENABLED).

This is also where proposal 05-284-r3 - SPC-4: Self Describing Command Timeouts, defines how a device server returns the appropriate (suggested) task timeout value, that application clients should wait before aborting a task, or, using the QUERY TASK TMF Extended Response, querying the task's status before deciding whether to abort it or whether to give it another chance.

Problem

The problem is that the timeout value(s) returned by the target as described in 05-284r3 currently does not correspond to any task state. What then does such a timeout value mean ***in the context of SAM-4 and SPC-4?***

Analysis

A task executing at the target device and completing with the application client needs to be represented by *at least two* task states, TASK ENABLED and TASK ENDED, at the target. This implies that the timeout value specifies time spent in at least those two states.

Suggested changes

The suggested changes are that the timeout value specify time spent in the TASK ENABLED and TASK ENDED task state. This of course implies that TASK ENABLED and TASK ENDED state is bounded in time. This is not a problem since the time bounds are not specified in the Architecture. As far as the Architecture is concerned a time bound of 100 million years is just as good as 100 microseconds, for example.

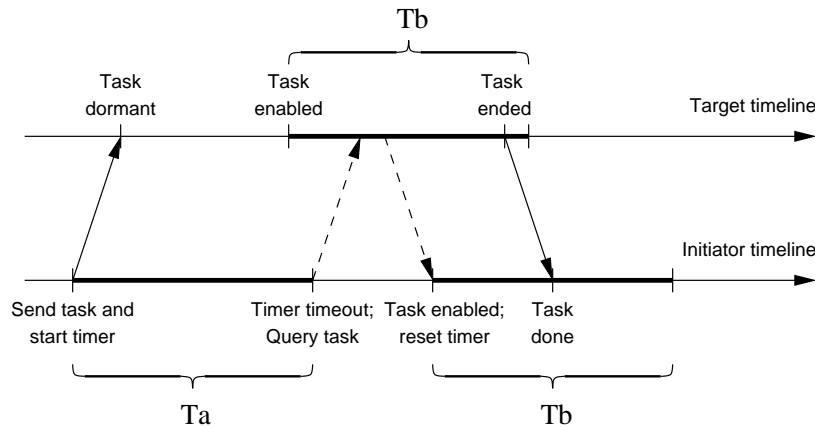
| Suggested changes are that this proposal, in a form as agreed to by the T10 committee, be included in SAM-4.

Implications

This would allow application clients to reset their timers if after querying the task, they receive a QUERY TASK TMF Extended response of TASK ENABLED or TASK ENDED, giving the task time to complete and avoiding error recovery "thrashing".

Since application clients do not detect task state transitions, the task timeout timer is already running at the initiator when the task arrives at the target.

When the timer times out and the task hasn't completed, the application client issues QUERY TASK TMF to query the task's location and state. If the target responds with QUERY TASK TMF Extended Response, see 06-368r0, and the task is in the TASK ENABLED or TASK ENDED task state then it would finish in at most the timeout value specified in 05-284r3, plus any transport overhead. The following figure shows an example of this.



Where,

T_a is a timeout value which is *not* specified.

T_b is a timeout value as specified in 05-284r3.

Depending on their environment, application clients may choose a smaller, equal or larger than T_b value for T_a . Application clients may also choose that T_a be equal or greater than the NOMINAL COMMAND PROCESSING TIME if defined, as specified in proposal 05-284r3. The value of T_a depends for example on the environment where the target device is, the SCSI Domain, usage patterns, whether the device is shared, number of initiators sharing the device, the protocol being used, etc. The value of T_a is not specified in the Architecture or in SPC-4.

It is recommended that they choose a value for T_a greater than or equal that of T_b to cover the most general cases of task processing. For example, a task with TASK SIMPLE or TASK ORDERED attribute, when there are no blocking or ordering conditions would become TASK ENABLED as soon as it enters the task set (no time is spent in TASK DORMANT state).

If the task is found to be in any other state, it is up to the application client's discretion to decide what to do next. (See [previous text in this document](#).) For example, if the task is found to be in TASK DORMANT state, the application client may wish to abort the task and notify the system administrator that logical unit needs external intervention.

Summary

This proposal gives a possible meaning to the timeout value(s) as specified in proposal 05-284r3 in the context of SAM-4 and SPC-4. It also gives a relationship between those timeout values and other task timeout values an application client may wish to use.

This proposal also outlines an error recovery framework to be used by application clients when communicating with targets implementing proposal 05-284r3 and proposal 06-368r0.

Note however, that targets are not required to implement either one of those two proposals. In such a case the old (current) behavior applies.

The intent of this proposal is that it be included in a form as agreed to by the T10 committee in SAM-4 along with 06-368r0, thus giving greater flexibility and more efficiency in error recovery mechanisms, while having little to no impact on the Architecture.