

**CONGRUENT SOFTWARE, INC.**  
**3998 Whittle Avenue**  
**Oakland, CA 94602**  
**(510) 531-5472**  
**(510) 531-2942 FAX**

FROM: Peter Johansson  
TO: NCITS T10  
DATE: February 28, 1998  
RE: SBP-2 Comment Resolution: Isochronous or Not?

---

This memorandum addresses a singular comment on SBP-2 Revision 3 submitted by Steve Finch. The text of the comment is reproduced below:

This document provides a definition of several basic functions: access control; task management; asynchronous command, data and status transportation; isochronous command, data and status transportation; isochronous control and status transportation; and, isochronous control commands. The definition and documentation of most of these are fairly mature with the exception of the those sections dealing with isochronous operation. This, I believe, is a direct result of insufficient participation by experts from the A/V and consumer electronics fields and a lack of attention to the subject matter by the storage manufacturers present. While we have made a valid attempt to obtain the inputs and participation from the A/V and consumer electronics industry, and some of the participants from storage manufacturers have spent a lot of individual time studying documents generated by other groups, I believe we can not assume we have accomplished this goal, especially not to the extent of setting it in concrete, i.e., including it into a standard. I recommend that the information pertaining to isochronous commands and isochronous operation be removed from this standard and that a new project be authorized to document these aspects. The need for isochronous storage devices is not so near that it needs to be documented immediately. I believe it is much more important that such documentation be complete and correct. I know of individuals who are willing to chair such a group and edit such a document. I have input from some A/V manufacturers that they want to pursue this mode of operation starting this year, so their participation can be expected. I think we had the best of intentions when we included isochronous operations in the scope of this project, but it was just too early to get the appropriate attention of those with the appropriate knowledge and experience. I do not want to see the asynchronous portion of this standard delayed.

The short response to the comment is that I believe that the acceptance and usage of isochronous SBP-2 devices will be advanced if the draft standard is forwarded as it stands and would be retarded if the isochronous material were removed. But a thoughtful comment deserves consideration and response to its particulars; these responses follow.

The pragmatic issue is not whether the material in SBP-2, as revised by responses to the T10 letter ballot comments, is fully matured. No standard, unless it is already moribund, can be fully matured. Do we assist future progress by establishing a milestone with the publication of SBP-2, inclusive of the isochronous material? I think we do and offer the following for consideration:

- The isochronous data recording format is widespread among nascent Serial Bus devices—even those that do not utilize the SBP-2 transport protocol. The SBP-2 draft standard is the only public document that specifies this recording format.
- Some particulars of the isochronous data recording format differ from the data formats specified by IEC FDIS/61883-1, the document from which the SBP-2 formats were extended. For example, SBP-2 describes how the CIP *fmt* field indicates whether or not the associated application data has embedded time stamps. Although the use of the most significant bit of *fmt* to differentiate the two data types is not in conflict with IEC FDIS/61883-1 and although the change embodied in SBP-2 was widely reviewed and accepted by the original architects of the CIP format, the change was never made in the IEC document.
- The two-queue model described by SBP-2 is fundamentally sound for all isochronous devices. In some cases, the second queue is implicit in the design of the device, as is the case when the device has embedded knowledge of its own file system; these devices can present a superficial appearance of implementing a single queue. The essence of my argument here is that the underlying architecture of isochronous SBP-2 rests on a solid foundation. We should not discard the work to date because of anxiety that some details may need work.
- Despite the lack of face-to-face participation in the *ad hoc* working group meetings, SBP-2 has received careful review by the consumer electronic industry. I have been part of some of those reviews and have exercised diligence to solicit comment and make corrections as appropriate.

There is precedent within T10 and its ancestor committees to move forward with standards work despite risk. The descriptions of synchronous parallel SCSI became part of a standard before any implementation efforts were underway. Their inclusion paved the way for the work that followed.

In light of the above arguments and the overwhelming approval of SBP-2 Revision 3 in T10 letter ballot, I recommend that the comment to remove the isochronous portions of SBP-2 be rejected.