

HEWLETT-PACKARD

Office Products Division

(619) 655-8553

FROM: Greg Shue
TO: NCITS T10
DATE: March 30, 1998
RE: Reconnect time-out

SBP-2 Revision 3d specifies a reconnect interval fixed at one second. While this may be sufficient for the resumption of access control subsequent to a Serial Bus reset where neither the initiator, target nor segments of the bus in between the two disappear from the topology, printer vendors believe it is too short in other circumstances. Consider the change caused when a user inadvertently disconnects a cable, unwittingly disconnects a cable between initiator and target or makes a new connection that results in a looped topology. Any of these transient interruptions may be quickly remedied by the user, but not within the current one second limit!

The cases described above will occur, since Serial Bus is a self-configuring, hot-pluggable bus that invites user changes. As currently specified by SBP-2, the target would be forced to cancel tasks and release resources. When the connection is restored the prior (incumbent) initiator may have already lost its access rights to another (newly) logged-in initiator. To remedy this situation, the following changes are recommended for SBP-2:

- Change the reconnect time-out to be a vendor-dependent parameter;
- Modify the LOGIN request ORB so that the initiator may specify the desired reconnect time-out and modify the login response data so that the target confirms the actual time-out value; and
- Provide some way for other initiators to determine how long resources will be held for initiator(s) that have yet to reconnect.

The text of these proposals follows; it references SBP-2 Revision 3d.

5.1.3.1 Login ORB

Before any other requests (except QUERY LOGINS) can be made of a target, the initiator shall first complete a login procedure that uses the ORB format shown below.

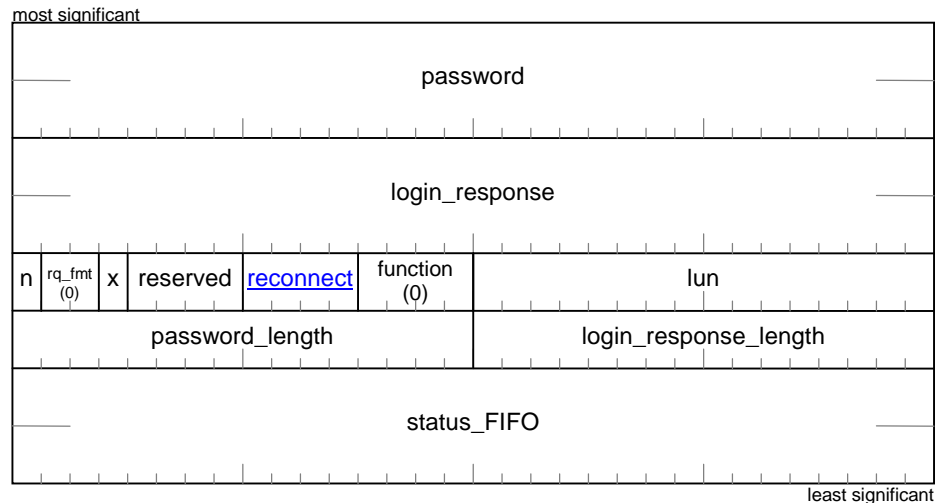


Figure 16 – Login ORB

The *password* and *password_length* fields may contain optional information used to validate the login request. If *password_length* is zero, the *password* field may contain immediate data. When *password_length* is nonzero, the *password* field shall conform to the format for address pointers specified by Figure 9 and shall contain the address of a buffer in the same node as the initiator; consequently the *node_ID* field of this address pointer is reserved. The buffer shall be accessible to a Serial Bus block read request with a data transfer length less than or equal to *password_length*. The format and usage of password data, whether immediate or indirectly addressed, are specified by Annex C.

The *login_response* and *login_response_length* fields specify the address and size of a buffer allocated for the return of the login response. The *login_response* field shall conform to the format for address pointers specified by Figure 9 and shall address the same node as the initiator; consequently the *node_ID* field of this address pointer is reserved. The buffer shall be accessible to a Serial Bus block write request with a data transfer length less than or equal to *login_response_length*. The initiator shall set *login_response_length* to a value of at least 12; the target may ignore this field [if it stores no more than 12 bytes of login response data](#).

The *notify* bit and the *rq_fmt* field are as previously defined for management ORB formats.

The *exclusive* bit (abbreviated as *x* in the figure above) shall specify target behavior with respect to concurrent login to a logical unit. When *exclusive* is zero, the target, subject to its own implementation capabilities, may permit more than one initiator to login to a logical unit. If *exclusive* is one the target shall permit only one login to a logical unit at a time; see 8.2 for a description of target behavior.

[The *reconnect* field shall specify the desired reconnect time-out as \$2^{\text{reconnect}}\$ seconds. The default reconnect time-out, when *reconnect* is zero, is one second. The target may not be able to support the requested value; see *reconnect_hold* in the login response data below.](#)

The *lun* field specifies the logical unit number (LUN) to which the request is addressed.

The *status_FIFO* field is as previously defined for management ORB formats and shall contain an address allocated for the return of status for the LOGIN request, status for all subsequent requests signaled to the *command_block_agent* allocated for this login and any unsolicited status generated by the logical unit.

If the login fails the contents of the response buffer are unspecified. Otherwise, upon successful completion of a login, the ~~response is returned in the format illustrated below~~ target shall store a minimum of 12 bytes of login response data and may store up to the entire 16 bytes illustrated below so long as the amount of data stored is an integral number of quadlets. Truncated login response data shall be interpreted as if the omitted fields had been stored as zeros.




Figure 17 – Login response

The *length* field shall contain the length, in bytes, of the login response data ~~and shall be equal to 12.~~

The initiator shall use the *login_ID* value returned by the target to identify all subsequent requests directed to the target's management agent that pertain to this login.

The *command_block_agent* field specifies the base address of the agent's CSR's, which are defined in 6.4. This field shall conform to the format for address pointers specified by Figure 9. The *node_ID* portion of the field shall have a value equal to the most significant 16 bits of the target's NODE_IDS register.

The *reconnect_hold* field shall specify the time, in seconds less one, that the target will hold resources for a previously logged-in initiator subsequent to a bus reset. The value of *reconnect_hold* shall not be greater than $2^{\text{reconnect}} - 1$, where *reconnect* is obtained from the login request. If an initiator fails to complete a successful reconnect request within *reconnect_hold* seconds after a bus reset, the target will perform a logout and release all resources held by that initiator (see 8.3).

5.1.3.2 Query logins ORB

The interpretation of the *login_ID* field in the query logins response data is modified so that other initiators may be aware of how long target resources are being held for initiator(s) that have not as yet reconnected. Replace the relevant paragraphs in this clause with the following:

The remainder of the query response is a variable-length array of 12-byte entries, one for each logged-in initiator, each of which contains a *node_ID*, *login_ID* and *initiator_EUI_64* field.

The *node_ID* field of an entry shall contain the node ID of a logged-in initiator. If a Serial bus reset occurred since the login was established and the initiator has not reconnected the login, the *node_ID* field shall have a value of FFFF_{16} .

NOTE – A *node_ID* value of FFFF_{16} may be observed only in the reconnect interval that exists for ~~one~~ *reconnect_hold* second(s) after a Serial Bus reset because after this time the target performs an automatic logout for any initiator that has not reconnected.

If the *node_ID* field has a value of FFFF_{16} , the *login_ID* field shall contain the time remaining, in seconds less one, until the initiator is automatically logged-out by the target. Otherwise, the *login_ID* field of an entry shall contain the login ID provided to the initiator as a result of its successful login.

The *initiator_EUI_64* field of an entry shall contain the EUI-64 obtained by the target from the initiator's configuration ROM at the time the login was validated.

7.4.n Reconnect Timeout entry

The Reconnect Timeout entry is an optional entry in the unit directory that describes the maximum reconnect timeout supported by a logical unit. Figure *nn* shows the format of this entry.



Figure *nn* – Reconnect Timeout entry format

3D₁₆ is the concatenation of *key_type* and *key_value* for the Reconnect Timeout entry.

The *max_reconnect_hold* field specifies the maximum value of *reconnect_hold* that the target may return in login response data (see 5.1.3.1). If this entry is not present in configuration ROM either the target does not include *reconnect_hold* in login response data or the value returned is always zero.

8.1 Access protocols

The first nonprocedural list item in this clause should be modified as shown below:

- The target implements a set of one or more *login_descriptors* that are used to hold context for logins. The context of a login stored in a *login_descriptor* consists of the *lun*, the *login_owner_ID*, the *login_owner_EUI_64*, the *status_FIFO* address, an *exclusive* variable, the base addresses of the fetch agent CSR's, the *login_ID* to be used by the initiator to identify the login, and the *reconnect_hold* period guaranteed by the target—these last ~~two~~three are returned to the initiator in the *login_response* data.

8.2 Login

The paragraph immediately below the nonprocedural list in this clause should be modified as shown below:

Once the above conditions have been met and a *login_descriptor* allocated, the initiator's *source_ID* is stored in *login_owner_ID*, the initiator's EUI-64 is stored in *login_owner_EUI_64*, the *lun* and *status_FIFO* fields from the login ORB are stored in the *login_descriptor*, the *exclusive* variable in the *login_descriptor* is set to the value of the *exclusive* bit from the login ORB, the addresses of the fetch agent(s) and the *reconnect_hold* value chosen by the target are stored in the *login_descriptor*. Lastly the target assigns a unique *login_ID* to this login and stores it in the *login_descriptor*.

8.3 Reconnection

The second paragraph in this clause should be modified as shown below:

~~For at least one second subsequent to a bus reset~~For each login, the target shall retain, for no longer than *reconnect_hold* second(s) subsequent to a bus reset, sufficient information to permit an initiator to reconnect ~~its~~that login . After ~~two~~*reconnect_hold* second(s) the target shall perform an implicit logout ~~for all login ID's that have~~if the login has not been successfully reconnected to ~~their~~its original initiator(s). The *reconnect_hold* parameter is communicated from the target to the initiator as part of the login response data.