

Virtual node Ids discussion nodes

From 7/15-16/98 meeting

- 1) Bridge portals [somehow] automatically assign ID to each device on local bus that "wants" one.
- 2) Portal has a map of virtual ID to PHY_ID. This must be updated on each bus reset (either by tracking topology or reading EUI-64s).
- 3) Virtual IDs stored in a requestor time out after t_3 time since last access from that requestor. (How does the requestor know if it is a physical ID or a virtual ID? Answer: The bus_id is \neq the requestor's own bus_id.)
- 4) Virtual IDs no longer used (because portal sees the device is no longer present) shall not be returned to the free pool until time t_4 since the last access. $t_4 > t_3$.
- 5) Virtual IDs on a bus must be the same as seen through all bridges attached to the bus.
- 6) Assume that mapping is done at the input of the outbound queue, and at the output of the inbound queue.
- 7) On bus reset, attached portals:
 - a) Freeze inbound async queues (packets addressed to that bus) (may discard, but not required)
 - b) Reexamine the topology of the reset bus (by means of the self-id packets).
 - i) Disconnected nodes – place the virtual node_id (if it had existed) into "limbo" (don't re-use until timeout t_4 expires, reject inbound packets with some error code)
 - ii) Newly connected nodes – assign them a new virtual node_id (as soon as one is available)
 - c) If a new bridge is connected, synchronize virtual IDs
 - d) Thaw the inbound queues
- 8) (not specific to virtual node-ids: Assuming t-labels are re-used by requestor immediately subsequent to a bus reset, then the portals on the reset bus must quarantine the requestor on the reset bus (for twice the split timeout) to prevent sending new requests and acceptance of a response from another bus that duplicates a t-label.)
- 9) The portal immediately adjacent to the packet originator virtualizes the source_id.
- 10) The portal immediately adjacent to the packet destination un-virtualizes the destination_id.
- 11) TBD: Evaluate all cases of connecting two bus segments for correctness.
- 12) TBD: Method to establish & synchronize virtual_ids (initialization).

From 8/20-21/98 meeting

For a local txn if a reset occurs after ack-pending & before response, the responder must discard the txn processing (& must not respond.) For a remote txn, if device is bridge aware then the responder can continue to respond even if bus reset.

(In SBP2 example) legacy requester is OK due to quarantine ideas.

Keep in mind AV/C: both devices (src & dest) will be txn initiators (for command & for status.

If an async packet (req or resp) has 3FF in src bus_id field of packet being routed off-bus, then bridge should map this to the real bus_id. ??????

Remember in our designs to deal with net topology changes (such as connecting 2 buses that are parts of 2 sep networks, or split net into two.

How should portals on one bus agree on virtual id mapping?

- 1) if self-id packets have bridge bit, then after self-id process the alpha portal is the highest phy-id bridge portal – if phy-ids change due to bus reset, then which portal has the well-known virtual node-id may change – so (due to principles of vnids) this node cannot be visible externally for a while
 - a) so keep in mind that even if we have alpha-portal idea, it may not be available to nodes off-bus at times
 - b) however, it IS instantly available on-bus
 - c) a newly alpha portal checks to see if it has a mapping table
 - d) if it doesn't, it builds one
 - e) if it does, it at least must re-verify it
 - f) during this rescan of the bus, see if limbo-ing an old vnid or some device wants a vnid then issue one (if available)
 - g) if new device wants a vnid but none is currently available, then during the unavailable time attempts by that device to communicate off-bus will get ack_XXXXXXXXXXXXX ????
- 2) note that in SBP2 case that the controller has to deal with the vnid timeout situation – how????
 - a) there was some discussion about using a wrapper on packets forwarded between bridges, and the bridge on the responder's bus uses a time stamp in the wrapper to determine that the vnid used in the packet may be stale