

802.1Qbg VSI Discovery and Configuration State Machine

Paul Bottorff

Paul Congdon

V12

August 17, 2010

Agenda

- Tighten up the VDP state machine specification for the 802.1Qbg draft
- Retain existing VDP protocol

Local Change Event Conventions

- Current VDP state machines uses the shorthand term “localChange-” to indicate a local event from the Station or the Bridge control plane.
- Other 802.1 state machines use variable tests to identify events.
- Definitions to identify events at Station:
 - sysCmd = request TLV from the Station’s system
 - rxResp = most recent Bridge response TLV received at the Station
 - A new command event is indicated by sysCmd != NULL
 - A new response event is indicated by rxResp != NULL
- Definitions to identify events at Bridge:
 - rxCmd = most recent Station command TLV received at the Bridge
 - Resp = response TLV constructed by the Bridge resourcing module
 - bridgeExit = a boolean which is set TRUE whenever the Bridge wants to exit
 - A new command event is indicated by rxCmd != NULL
 - A Bridge resource response is indicated by Resp != NULL
 - An Bridge exit event is indicated when bridgeExit = TRUE

Command-Response TLV Field Conventions

- All TLVs have the same format as in draft 1.0
- Mode1 is 1 octet encoding the command or response
 - PREASSOC: 0
 - PREASSOC (with resources): 1
 - ASSOC: 2
 - DEASSOC: 3
 - Note: Response does not always reflect command (DEASSOC case)
- Mode2 is 1 octet encoding the completion code
 - SUCCESS: 0
 - Fail Codes: 1-255
- State machine may references TLV fields as follows:
 - operCmd TLV: operCmd.Mode1, operCmd.Mode2 and operCmd.ExceptMode2
 - rxCmd TLV: rxCmd.Mode1, rxCmd.Mode2 and rxCmd.ExceptMode2
 - Resp TLV: Resp.Mode1, Resp.Mode2 and Resp.ExceptMode2
 - rxResp TLV: rxResp.Mode1, rxResp.Mode2 and rxResp.ExceptMode2
 - Note: TLV.Mode1, TLV.Mode2 and TLV.ExceptMode2 = NULL when TLV = NULL
 - Note: TLV.ExceptMode2 refers to all fields of the TLV except Mode2
 - Note: Current state machine uses shorthand for the content of the TLV mode field, ASSOC_ACK_rx, ASSOC_NACK_rx, PREASSOC_ACK_rx, PREASSOC_NAK_rx, DEASSOC_ACK_rx

Timer Events Conventions

- Current Station state machine uses two timers for a response timeout and for a keep alive.
- Current Bridge state machine uses one timer for a keep alive time out.
- Change Station to a single count down timer
 - StartRespTimer() to `waitWhile` = `respWaitDelay`
 - `ACTIVITY_TIMER_EVENT` to `waitWhile` = `reinitKeepAlive`
 - Response timeout when `waitWhile` == 0
 - Keep alive event cycle when `waitWhile` == 0
- Change Bridge to a single count down timer
 - `INACTIVE` to `waitWhile` = `toutKeepAlive`
 - Keep alive timeout when `waitWhile` == 0
 - Add `waitWhile` = `resourceWaitDelay`
 - Response timeout when `waitWhile` == 0

Expressing Bridge Resourcing States

- Current state machine expresses the Bridge side resourcing request as a procedure hand-off.
- The bridge state machine hands the command to another machine which may query network management system before sending a response.
- It is convenient to break the command request into two states to allow time-out for the command hand-off
 - Need a timeout for the resource request from the Bridge
 - Need to build the response for the station from the resource request response

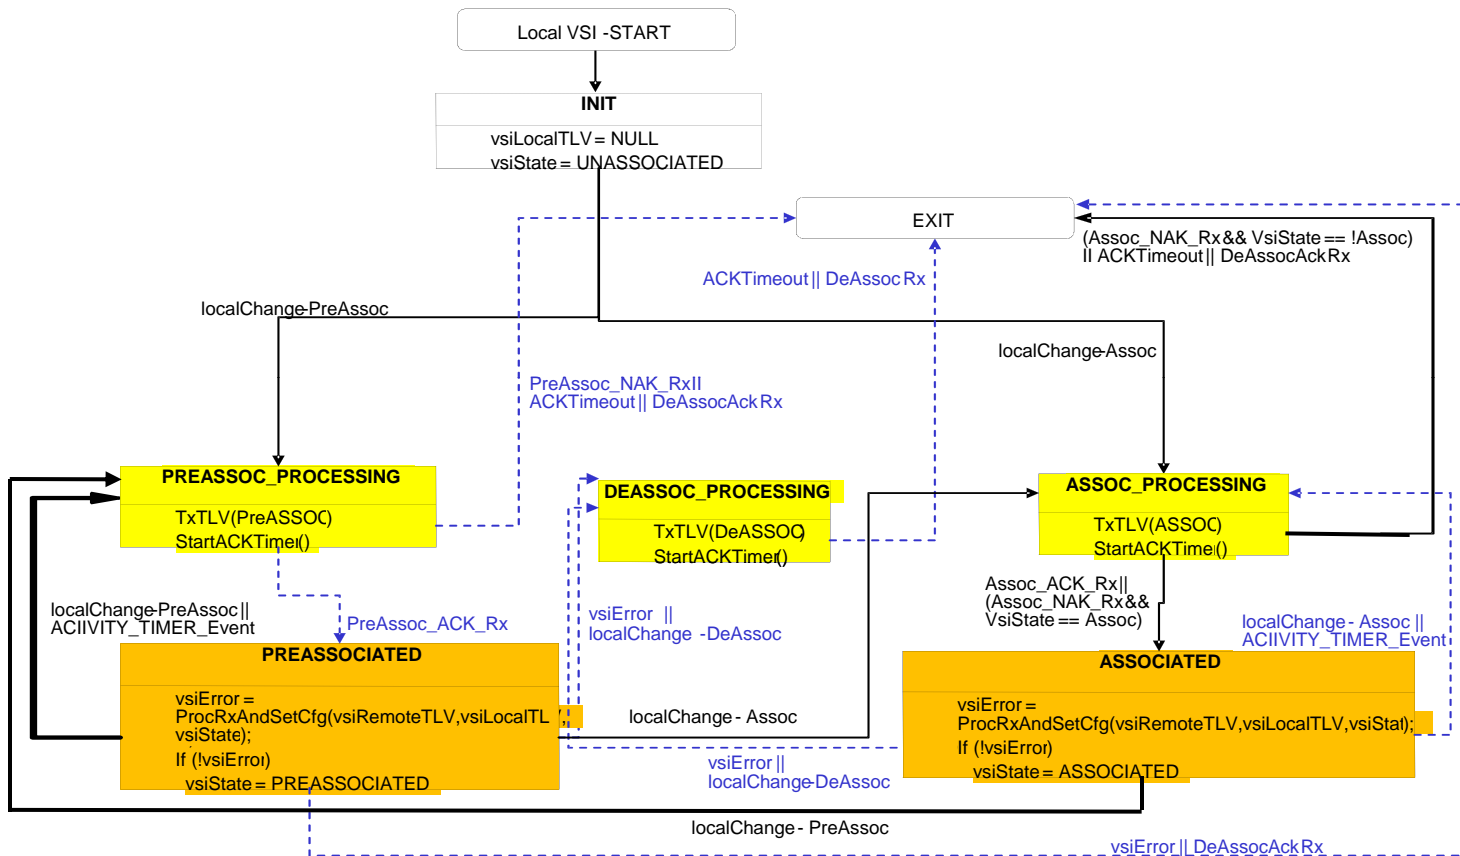
Expressing Reverting to Last Successful Associate

- The current VDP machines reverts to the previous successful associate in the event of an associate failure.
- The machine could be more specific about retaining the existing state.
- Also the machine needs to improve the clarity of handling a failed keep alive.

Error Processing and if...then... elimination

- Current VDP machine illustrates error checks on procedure calls, however these are not conventionally illustrated in our machines unless we specify the specific errors possible and how they affect the machine outcomes.
- We suggest removing these from the illustration unless we fully specify the operation of the procedures.
- Also we normally want to remove the if...then...s from state blocks to fully expand the state machine.

VDP Draft 1.0: State consolidation



- Yellow states may be consolidated
- Orange states may be consolidated

Procedures in Proposed VDP Machine

- TxTLV(TLV) sends a TLV to the VDP listener for packing into a ECP frame. The VDP listener also receives all VDP TLVs from ECP, unpacks the TLVs, may create a new VDP bridge machine if none already exists, and places the received TLVs in the variables rxResp on the station and rxCmd on the Bridge. On the station rxResp is filled with the current TLV whether or not it already has a TLV. On the bridge the VDP listener must wait for rxCmd to be NULL after creating a new VDP listener before filling rxCmd.
- Bridge Procedures:
 - resourceCmd(rxCmd) makes a resource request from the Bridge and builds a response TLV for the command placing it in the global variable Resp. The response may be PREASSOC, ASSOC or DEASSOC with success or fail. The response is returned in the variable Resp.
 - resourceCmd decodes each possible command type
 - DEASSOC: Response is DEASSOC with success. The Bridge machine moves to END.
 - PREASSOC: Response is PREASSOC with success if the resources requested are valid and fail if the resources are invalid. Otherwise the Bridge responds with fail or with DEASSOC and moves to END.
 - PREASSOCR: Response is PREASSOC with success if the resources requested are valid, are available and reserved for this VSI. Otherwise the Bridge responds with fail or with DEASSOC and moves to END.
 - ASSOC: Response is ASSOC with success if the resources requested are valid, are available, and enabled. Response is fail if the resources are invalid, unavailable, or not enabled. If fail the Bridge will move back to operating ASSOC or to END if no ASSOC was operating at the time of the request. The resourceCmd response may be ASSOC success, ASSOC fail, or DEASSOC success. A DEASSOC success indicate the Bridge will not continue while ASSOC fail indicates the Bridge will revert to the previous successful ASSOC.
 - resourceFree() frees all the resources associated with this state machine instance.
 - Dea = buildDea() builds a DEASSOCIATE TLV for the current state machine and returns it in the return parameter.

Timers

- Station timers
 - `waitWhile` = `respWaitDelay`
 - `waitWhile` = `reinitKeepAlive`
- Bridge timers
 - `waitWhile` = `resourceWaitDelay`
 - `waitWhile` = `toutKeepAlive`

Variables

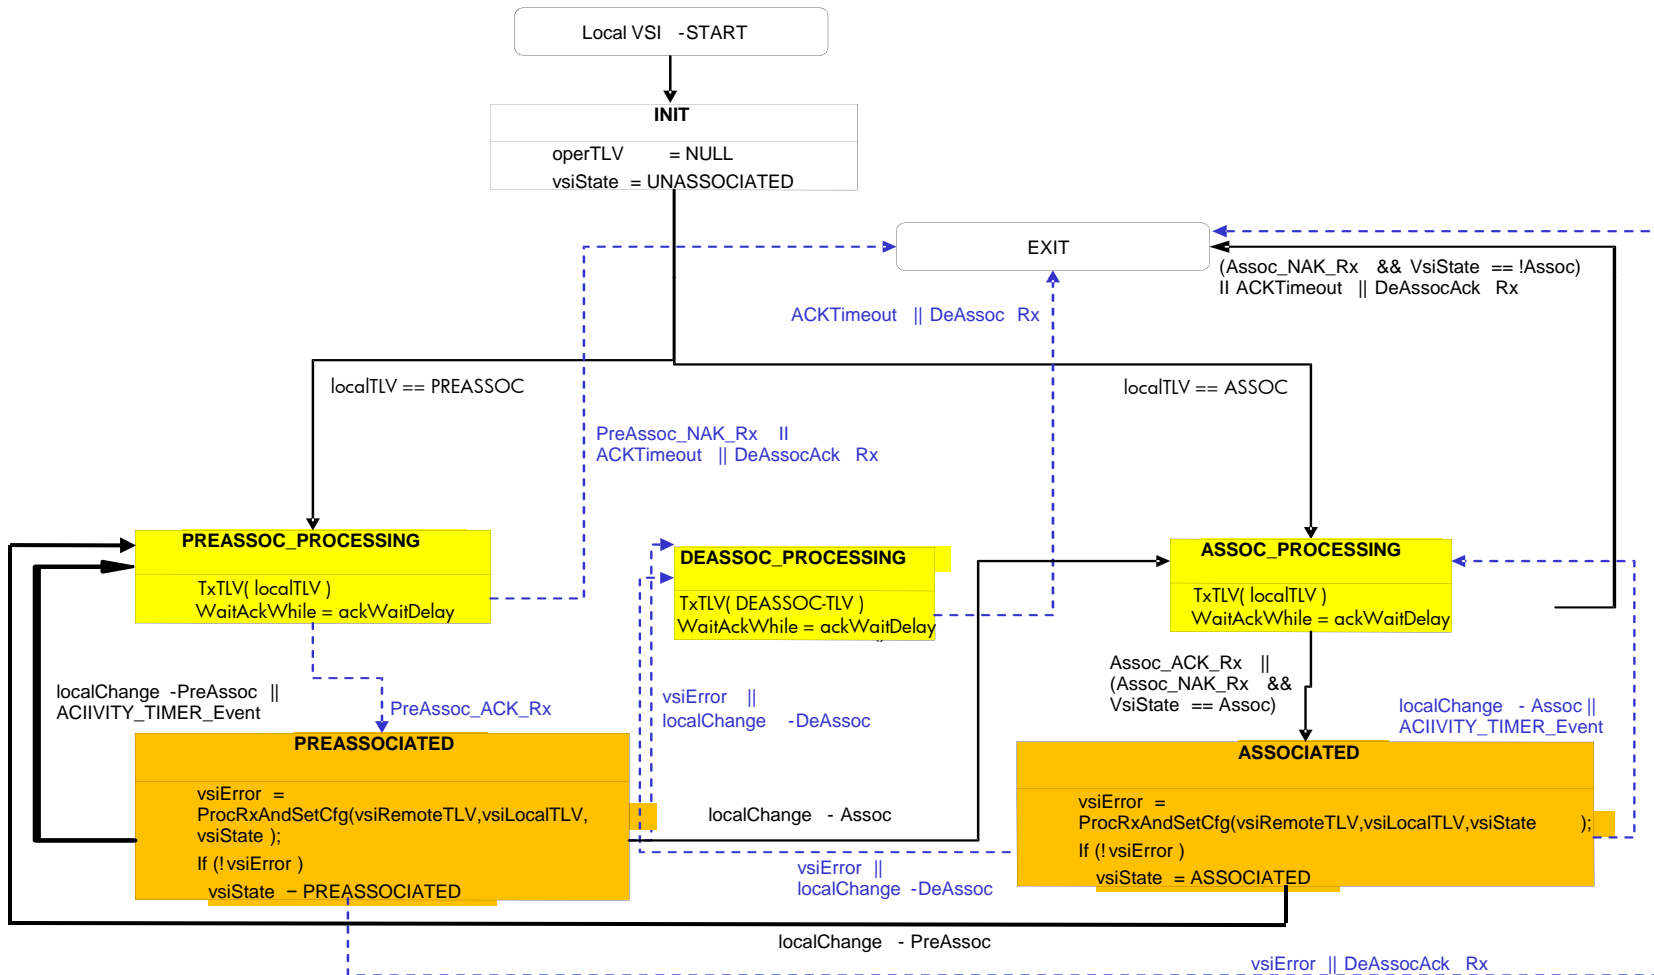
- sysCmd: A command TLV from the system or hypervisor or a NULL if no pending command. The VDP machine is ready to accept a new command when sysCmd is NULL.
- operCmd: The operating command TLV at the station or bridge or NULL if there is no operating command.
- rxCmd: The last received command TLV or NULL.
- Resp: A response TLV returned from the procedure resourceCmd. The variable may be set to NULL before issuing resourceCmd. The variable will be non-NULL when resourceCmd completes.
- rxResp: The last received response TLV. The rxResp contains the last received Resp TLV at the station or NULL. RxResp is NULL if no TLVs have been received or if cleared by the VDP machine. Note: it is possible to have a race when clearing rxResp since it may be updated asynchronously if the Bridge issues an unsolicited DEASSOC.
- Dea: A DEASSOC TLV constructed for this VDP machine
- vsiState: The current state of the VDP machine. This variable may take the values UNASSOC, PREASSOC, PREASSOCR, or ASSOC.
- bridgeExit: A boolean signal from the Bridge indicating the bridge VDP machine should exist. When TRUE the VDP machine should exit.

Summary

- New equivalent state machine provides more explicit state machine representation and consistency with IEEE practice.

BACKUP SLIDES

Station VDP State Machine: Draft 1.0



Bridge VDP State Machine: Draft 1.0

