

VSI Discovery and Configuration Protocol (VDP) (Working Draft)

*Proposed Resolutions to subset of comments against
802.1Qbg Draft-1*

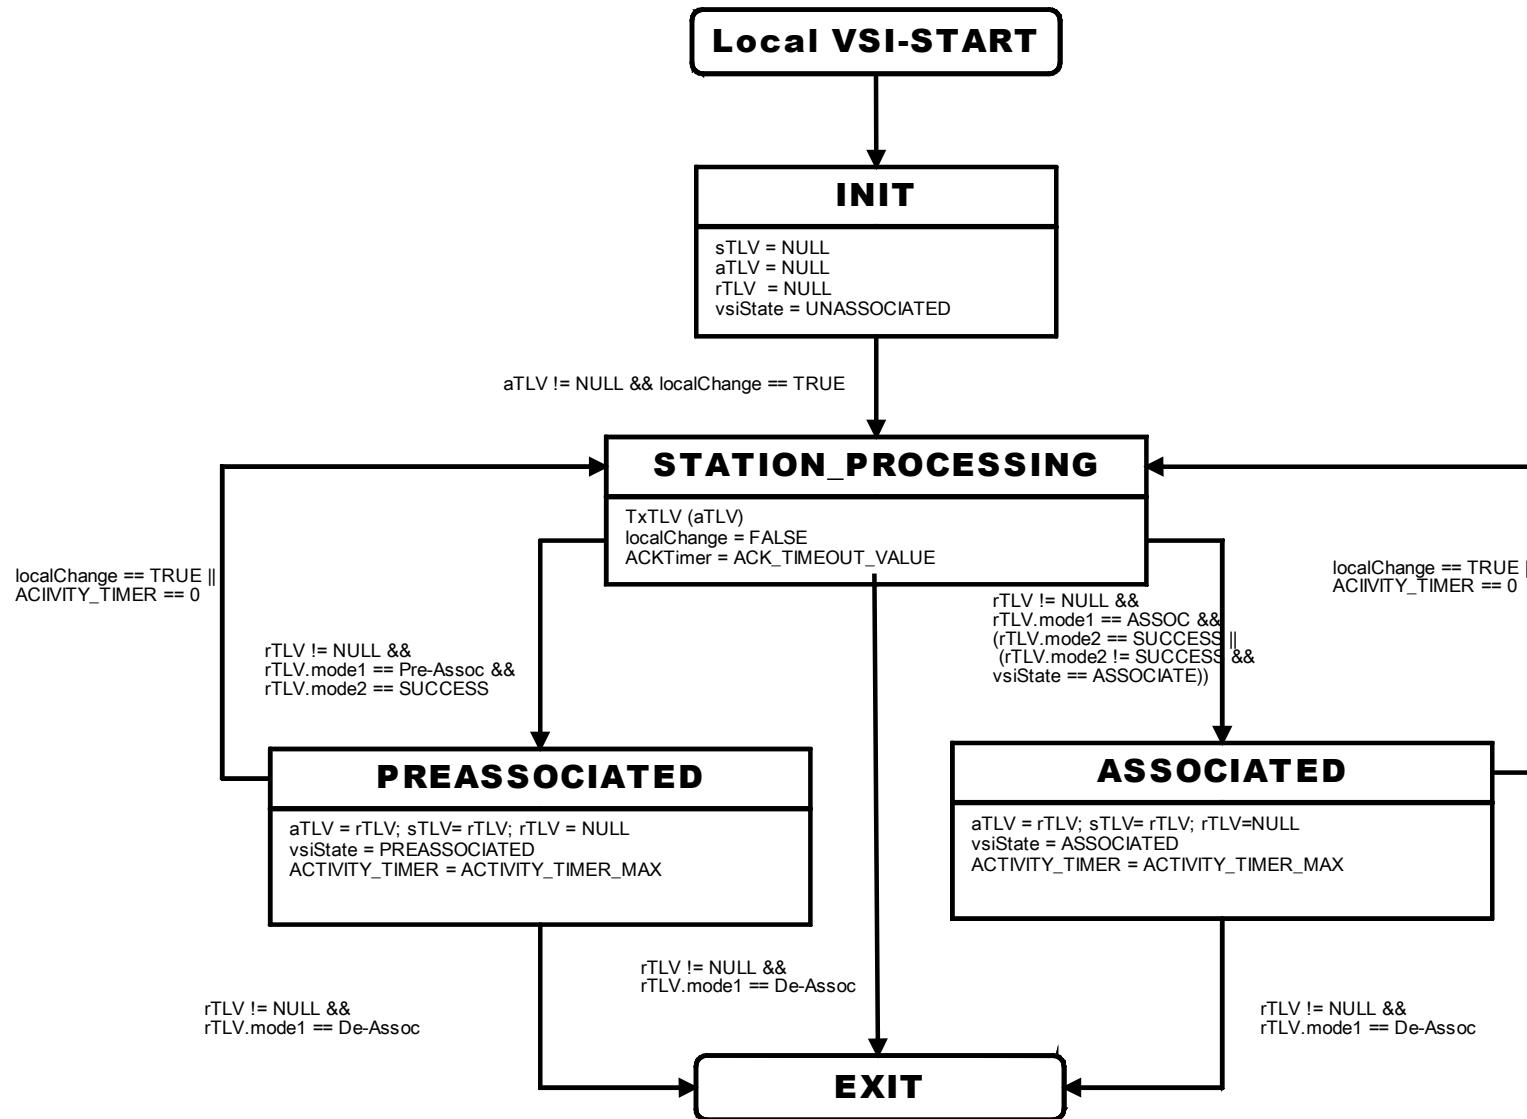
*Chait Tumuluri (Emulex), Jeffrey Lynch (IBM), Vijoy Pandey (BNT), Rakesh Sharma (BM),
Renato Recio (IBM), Srikanth Kilaru (Juniper)*

*Version 04
08/03/2010*

Summary of changes

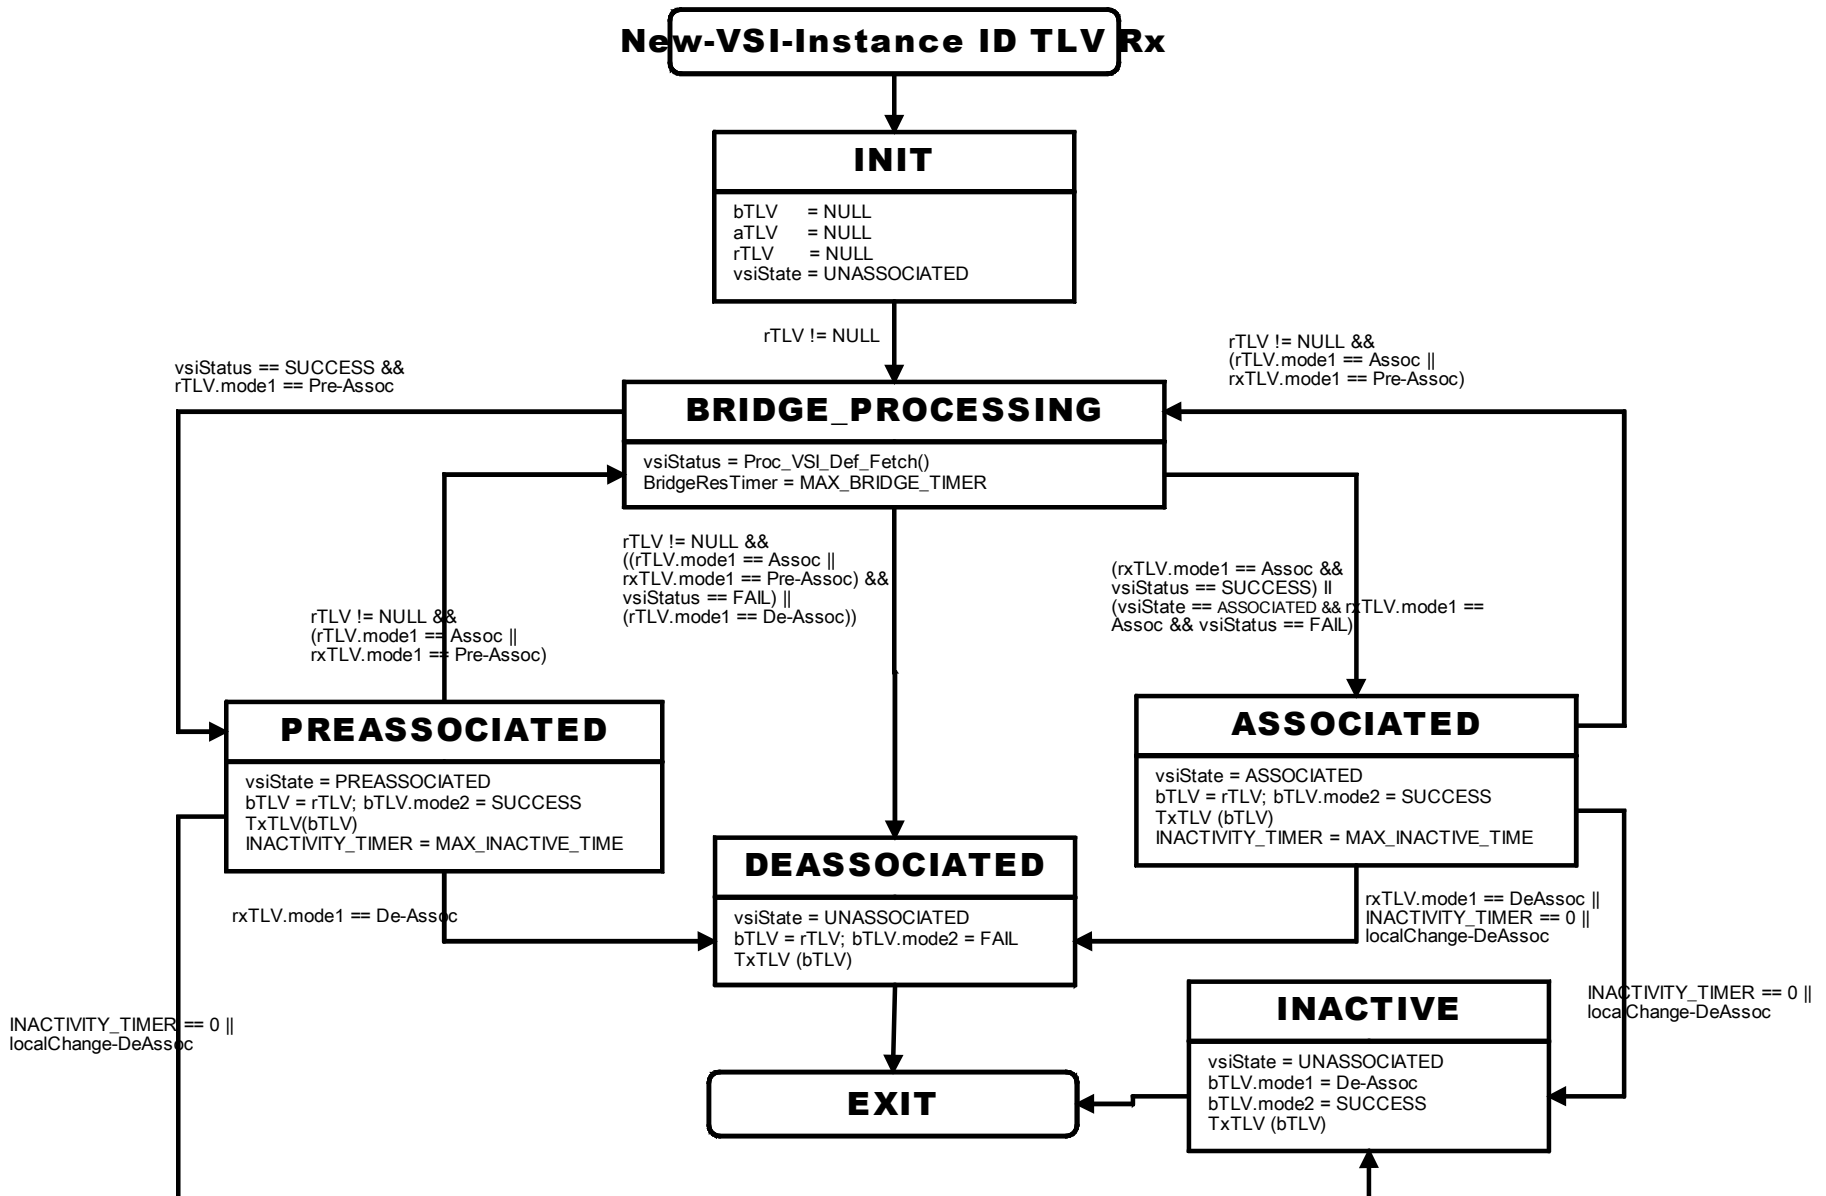
1. Compliance to IEEE Protocols State Machines Conventions (Ref: IEEE 802.1AB x.xx)
2. Changed timer expiration events to comply with IEEE state machines style.
3. Changed state machine layout to match IEEE state machines.
4. Misc. other style changes in state machines to make them consistent with IEEE state machines.
5. Updated local variables and procedures names and description for additional clarification.
6. ***Ensured that changes do not impact state machines operation other than style and layout.***

Station State Machine (One Instance per VSI)



VSI State Machine – Bridge (Draft)

(One Instance per VSI)



VSI State Machine Local Variables and Procedures

1. **vsiState:**

- Local variable for current state. Possible states are: UNASSOCIATED, PREASSOCIATED and ASSOCIATED

2. **sTLV:**

- Current local (active) TLV (configuration) of Station.

3. **aTLV:**

- TLV from local administration. In addition localChange variable is set.

4. **rTLV:**

- TLV received from remote (station or bridge)

5. **bTLV:**

- TLV received from remote (station or bridge)

6. **TxTLV(*vsiTLV*):**

- Transmits TLV using TLV transport (ECP) service interfaces. Includes support for aggregation of VSI TLVs.

7. **Proc_VSI_Def_Fetch ():**

- Bridge procedure to fetch VSI definition based on Associate and/or Pre-Associate parameters based on received Remote TLV parameters. And sets vsiStatus in case of error. This function handles PreAssociate with and without resource reservation case as well.

8. **ACKtimer:**

- ACKtimer is ACK Timer local variable. Response (ACK or NACK) is expected before timer expires.

9. **ACK_TIMEOUT_VALUE:**

- This local variable contains ACK timeout value for the machine.

VSI State Machine Local Variables and Procedures

1. **INACTIVITY_TIMER**

- Bridge timer local variable for inactivity timer.

2. **MAX_INACTIVE_TIME**

- Max inactivity time on the bridge. It is set by bridge based on configured value and complies to range specified in 802.1Qbg.

3. **BridgeResTimer**

- Bridge resource fetch timer local timer.

4. **MAX_BRIDGE_TIMER**

- Local variable that stores max bridge timer value. It is set by bridge based on configured value and complies to range specified in 802.1Qbg.

Backup

VDP Objectives

1. Support VSI preAssociate (with and without resource reservations), Associate and deAssociate.
2. ASSOCIATE, PreAssociate and DeAssociate are Idempotent i.e. can be repeated.
3. Capability to Associate skipping PreAssociate.
4. VDP will work both for VEPA and VEB environments.
5. Timeout mechanism to ensure:
 - a. Bridge resources are not reserved too long for inactive VSIs (lease semantics)
 - b. Allow removing resources from inactive VSIs with the goal of
 - a. Conserve bridges resources (Number VSIs being handled by bridge can be large).
 - b. Prevent inactive or VMs in error state to continue to hold resources.
6. Timeout out values to be negotiated on per channel between station and bridge. One timeout used for all ULPs on the channel negotiated using EVB TLV.

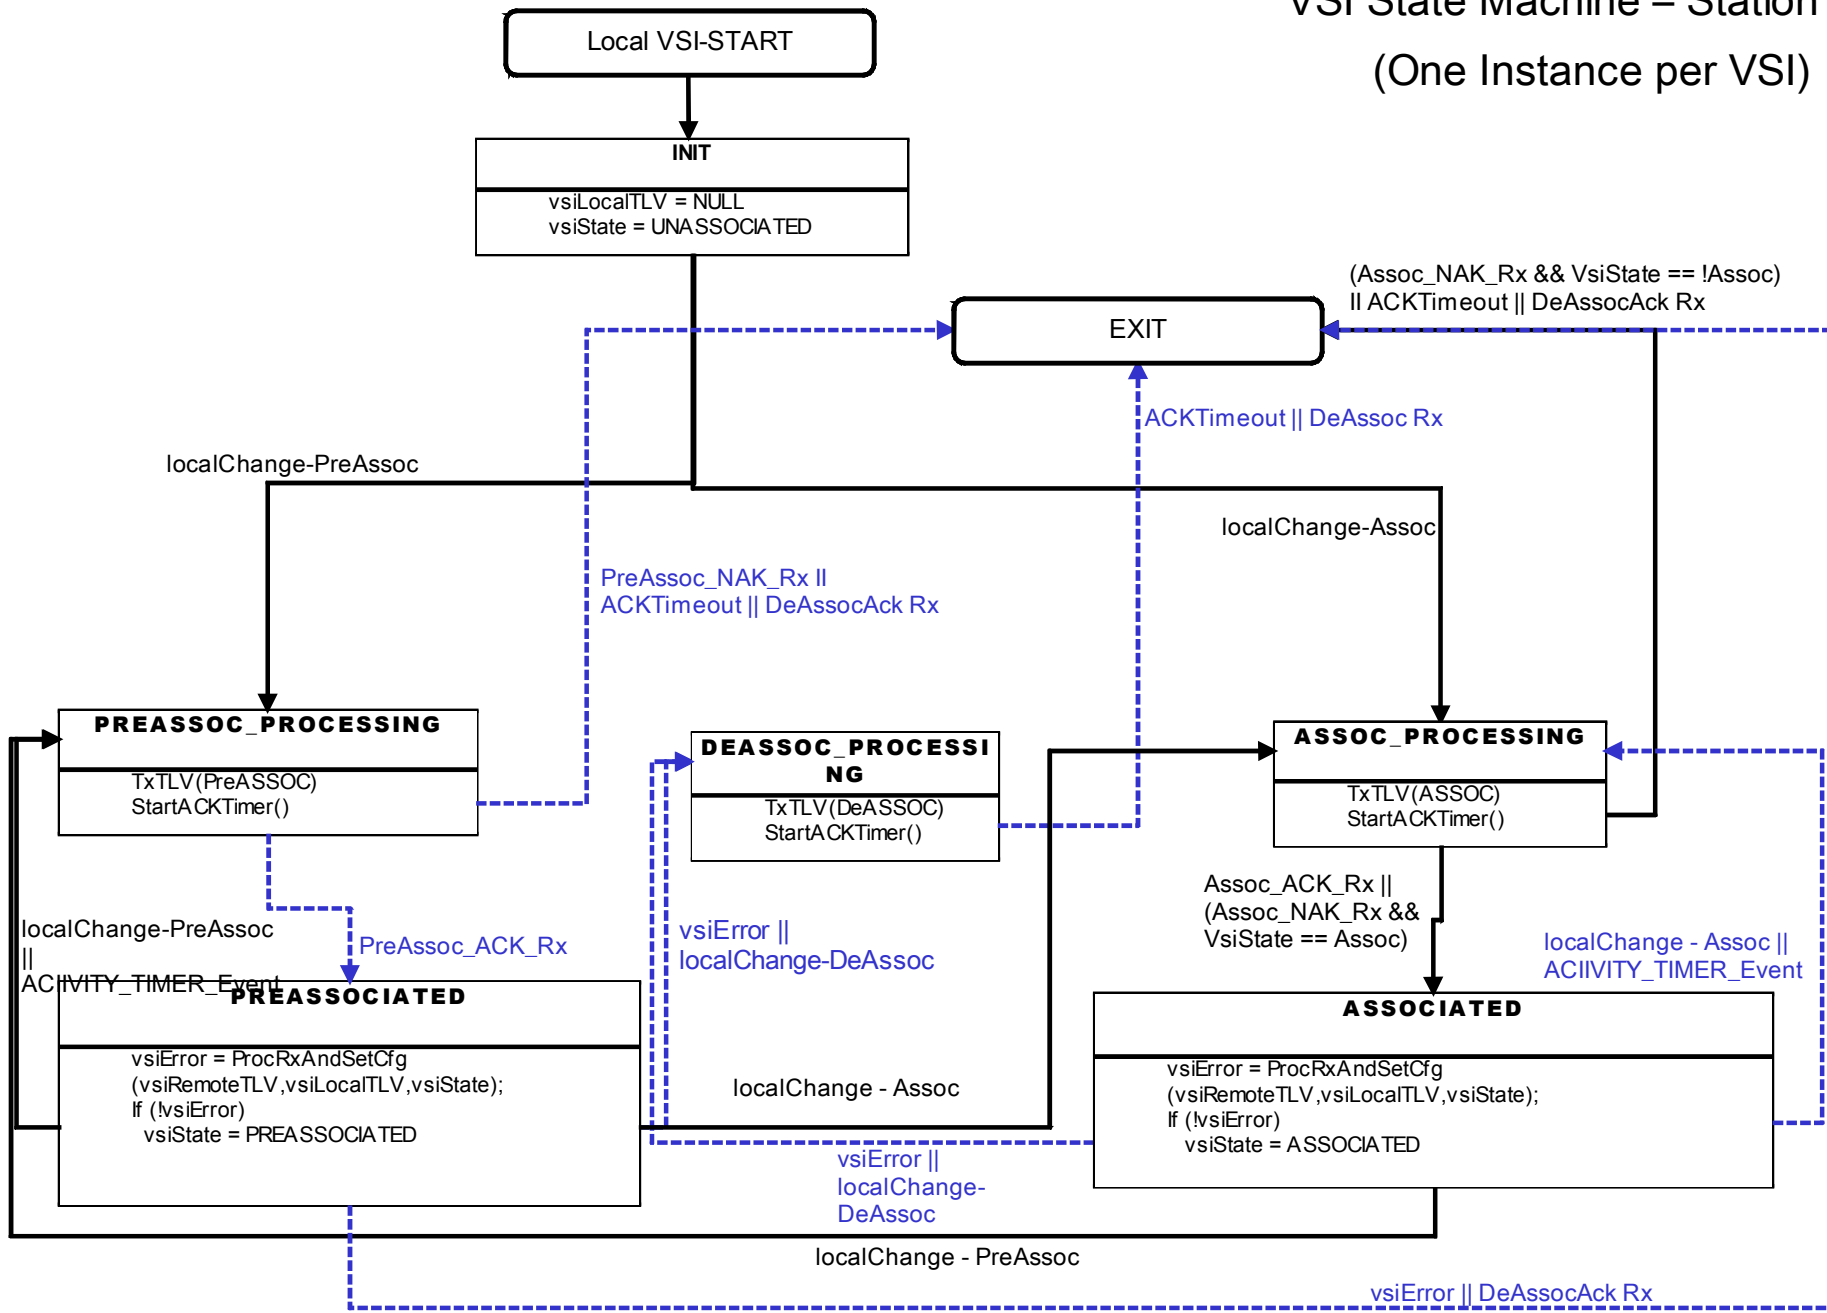
VDP Objectives (continued)

7. Manageability and Robustness

- a. Ensure VSI state and configuration between the Station and the Bridge remains consistent.
- b. Hard errors at the Bridge or the Hypervisor that can impact individual VSI or Hypervisor/ Bridge as a whole. All VSI configuration goes away in this situation.
- c. Bridge and Station Errors are detected through one or more of the following mechanisms.
 - VSI KEEP-ALIVE (periodic transmission of VSI TLV from station and response from Bridge)
 - ACK Timer
 - Transport (ECP and LLDP) status indications.
- d. Bridge and Hypervisor administrative capability to force VSI deAssociate.
- e. ASSOCIATED VSI remains associated with existing configuration if a new associate is not successfully completed.
- f. Statistics and logging support (need specific proposal).

VSI State Machine – Station

(One Instance per VSI)



VSI State Machine – Bridge (Draft) (One Instance per VSI)

