DCBX Framework

Joe Pelissier

az-pelissier-dcbx-framework-0309

DCBX Framework Overview

- **The DCBX Framework provides three capabilities utilizing LLDP:**

  Informational Passing: the passing of information that is used in DCB applications or configuration

  e.g., passing the number of traffic classes supported by a bridge to an attached end station may enable the end station to optimize its buffer resources

  Symmetric Parameter Passing: the passing of a parameter from one port to its peer port with objective of both ports utilizing the same parameter

  e.g., Priority Flow Control configuration

  The “to” and “from” ports are administratively configured

  e.g., passing configuration to and end station from a bridge

  Asymmetric Parameter Passing: the passing of a parameter from one port to its peer port, or between both peer ports. In this case, the resulting configuration of each port may or may not match its peer

  e.g., ETS bandwidth configuration

  The directions in which the parameters are passed is administratively configured

- **The Framework provides standard state machines for each of the parameter passing methods**

Asymmetric Parameter Passing Options

- **Two options for Asymmetric Parameter Passing have been discussed**
  - **Both are included in this presentation for comparison and (hopefully) drive consensus on one approach**
- **Option 1 optimizes simplicity and reduced number of transfers over octets consumed in the TLVs**
  - **In addition, the Option 1 state machine is nearly identical to the Symmetric Parameter Passing state machine**
- **Option 2 optimizes reduces octets consumed in the TLVs at the expense of a more complex state machine and more transfers required for convergence**

Asymmetric Parameter Passing – Overview (Option 1)

- **Two TLVs:**
  - **Configuration TLV: Provides current operational state and Willing bit**
    - Willing bit is not used by state machine
    - Provided to higher layers to provide an indication of the expected behavior of the remote port
  - **Recommendation TLV: Provides recommendation for the operational state of remote port**
    - Transmitted only if the local port is configured to make recommendations (in which case it is transmitted in all LLDP PDUs
    - Transmitted regardless of the “willingness” of the remote port
    - Recommendation Valid (RV) bit provides a mechanism for the local port to delay making recommendations until it receives additional information about the remote port (e.g. number of traffic classes supported and priority to priority group assignments).
- **Note: This same principle could be used with a single TLV that carries both the Configuration and Recommendation parameters**
  - **Two TLVs are proposed to optimize the case in which a recommendation is not being made, i.e. the recommendation TLV is not sent**

Asymmetric Parameter Passing State Variables (Option 1)

LocalWilling: Indicates that the local port has been administratively configured to accept recommendations. This value is included in the Willing bit of Configuration TLVs transmitted by this port.

OperParam: The current operational state of the parameter on the local port. This value is included as the parameter in the Configuration TLVs.

LocalAdminParam: The administratively configured state for the parameter. This becomes the operational state of the parameter by default, and may be overridden if the local port accepts a recommendation from the remote port.

RV: TRUE indicates that the RV (Recommendation Valid) bit was set in the last Recommendation TLV received. FALSE indicates that the RV bit was not set in the last Recommendation TLV received. NULL indicates either the TLV was not present in the last LLDP PDU or that no LLDP PDUs have been received.

Asymmetric Parameter Passing – Example TLVs (Option 1)

TLV Type =127	TLV Info String Len=17	802.1 OUI 00-80-C2	802.1 Subtype = 9	Reserved	Willing	Priority Assignment Table	Priority Group Configured Bandwidth Table
7 bits	9 bits	3 octets	1 octet	7 bits	1 bit	4 Octets	8 Octets

Configuration TLV

TLV Type =127	TLV Info String Len=13	802.1 OUI 00-80-C2	802.1 Subtype = 10	Reserved	RV	Recommended Priority Group Bandwidth Table
7 bits	9 bits	3 octets	1 octet	7 bits	1 bit	8 Octets

Recommendation TLV

Asymmetric Parameter Passing State Machine (Option 1)

Note: RV can take on three values: TRUE, FALSE, and NULL. NULL indicates that the Recommendation TLV was not part of the last LLDP PDU received, or that no LLDP PDUs have been received.

Asymmetric Parameter Passing Ladder Diagram (Option 1)

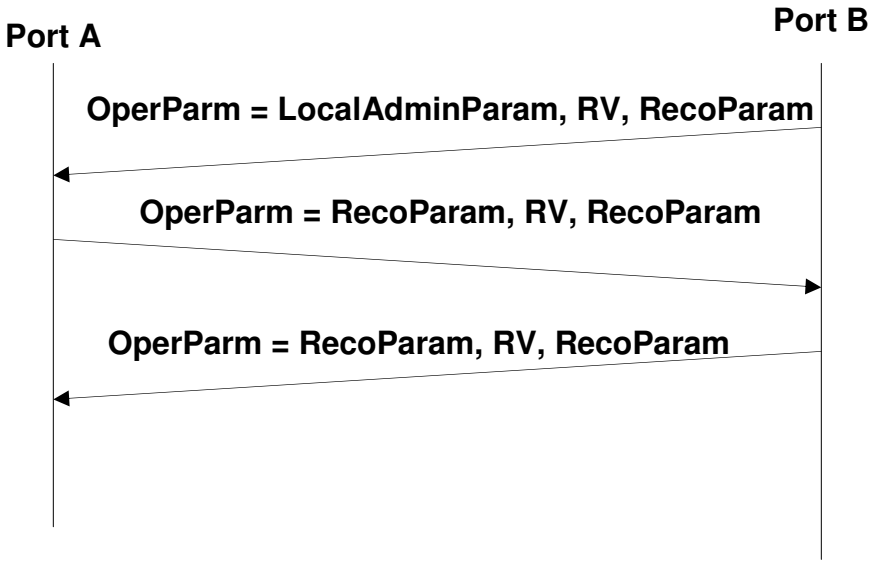

Two transfers in one direction and one transfer in the other direction required to transfer the parameters

Asymmetric Parameter Passing – Overview (Option 2)

- **One TLV:**

  The "parameter" field changes values depending on the state of the transmit state machine:

  Contains a recommended value when the state machine is actively making a recommendation

  Contains the current operational value otherwise

  Two state machines (transmit and receive) and a number of flags provide handshake for transfer of parameters

Asymmetric Parameter Passing State Variables (Option 2)

LocalReco: Indicates that the local port has been administratively configured to offer recommendations. This value is included in the Reco bit of TLVs transmitted by this port.

RemoteReco: TRUE indicates that the Reco bit was set in the last TLV received. FALSE indicates that the Reco bit was not set in the last TLV received. NULL indicates either the TLV was not present in the last LLDP PDU or that no LLDP PDUs have been received.

LocalWilling: Indicates that the local port has been administratively configured to accept recommendations. This value is included in the Willing bit of TLVs transmitted by this port.

RemoteWilling: TRUE indicates that the Willing bit was set in the last TLV received. FALSE indicates that the Willing bit was not set in the last TLV received. NULL indicates either the TLV was not present in the last LLDP PDU or that no LLDP PDUs have been received.

LocalRecoPresent: TRUE indicates that a recommendation is occupying the parameter field of the TLV. FALSE indicates that the current operational state is occupying the parameter field of the TLV. This value is included in the RecoPresent bit of TLVs transmitted by this port.

RemoteRecoPresent: TRUE indicates that the RecoPresent bit was set in the last TLV received. FALSE indicates that the RecoPresent bit was not set in the last TLV received. NULL indicates either the TLV was not present in the last LLDP PDU or that no LLDP PDUs have been received.

Asymmetric Parameter Passing State Variables (Option 2)

LocalRecoReceived: TRUE indicates that the local port is acknowledging the receipt of a recommendation from the remote port. FALSE indicates that the local port is not currently acknowledging a recommendation from the remote port. This value is included in the RecoReceived bit of the TLVs transmitted by this port.

RemoteRecoReceived: TRUE indicates that the RecoReceived bit was set in the last TLV received. FALSE indicates that the RecoReceived bit was not set in the last TLV received. NULL indicates either the TLV was not present in the last LLDP PDU or that no LLDP PDUs have been received.

LocalParam: The parameter to be transmitted in TLVs from this port. The parameters to be transmitted vary for each TLV and are identified in the normative clauses for each TLV. For example, in the ETS TLV, the parameter is the table of bandwidth assignments for each Priority Group.

RemotePraram: Contains the value of the last parameter received in the TLV (regardless of whether the parameter contained the operational parameter or recommended parameter, as indicated by the Reco bit). NULL indicates either the TLV was not present in the last LLDP PDU or that no LLDP PDUs have been received.

Asymmetric Parameter Passing State Variables (Option 2)

OperParam: The current operational state of the parameter on the local port. This value is included as the parameter in TLVs transmitted from this port when the Reco bit is set FALSE.

LocalAdminParam: The administratively configured state for the parameter. This becomes the operational state of the parameter by default, and may be overridden if the local port accepts a recommendation from the remote port.

LocalRecoParam: The current recommendation of the parameter setting that the local port wishes to make to the remote port. If the local port is recommending and the remote port is willing, this is the value that will become the operational state of the remote port. This value is included as the parameter in TLVs transmitted from this port when the Reco bit is set TRUE.

LocalAdminState: Indicates that the local port is operating off of its locally administered parameter (as opposed to parameters that were recommended by the remote port). Note that if as a result of this mechanism the local port accepts the remote port's recommendation, this value is set FALSE even if the recommended parameter matches the LocalAdminParam. This value is included in the AdminState bit of TLVs transmitted by this port.

RemoteAdminState: TRUE indicates that the AdminState bit was set in the last TLV received. FALSE indicates that the AdminState bit was not set in the last TLV received. NULL indicates either the TLV was not present in the last LLDP PDU or that no LLDP PDUs have been received.

Asymmetric Parameter Passing – Example TLV (Option 2)

TLV Type =127	TLV Info String Len=13	802.1 OUI 00-80-C2	802.1 Subtype = 10	AsymFlags	Priority Group Bandwidth Table
7 bits	9 bits	3 octets	1 octet	1 octet	8 Octets

AsymFlags:

Bits 7-5: Reserved
Bit 4: Willing
Bit 3: Reco
Bit 2: RecoReceived
Bit 1: AdminState
Bit 0: RecoPresent

Parameter: Priority Group Bandwidth Table

Asymmetric Parameter Passing State Machines (Option 2)

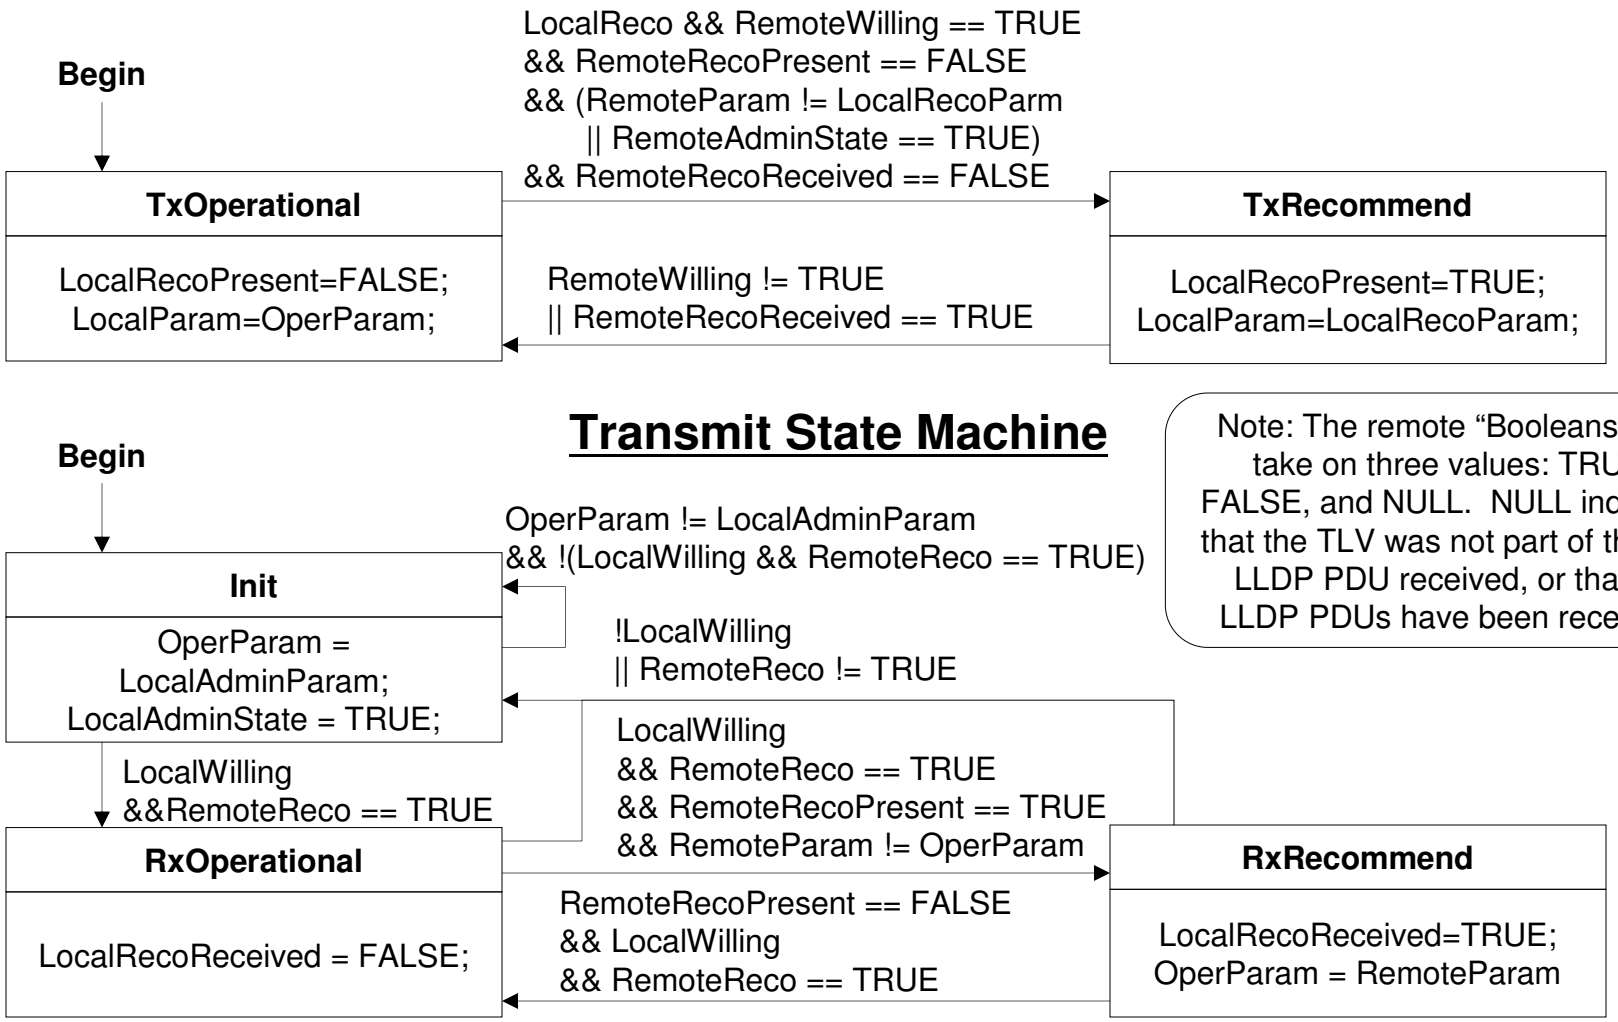

Note: The remote "Booleans" can take on three values: TRUE, FALSE, and NULL. NULL indicates that the TLV was not part of the last LLDP PDU received, or that no LLDP PDUs have been received.

Asymmetric Parameter Passing Ladder Diagram (Option 2)

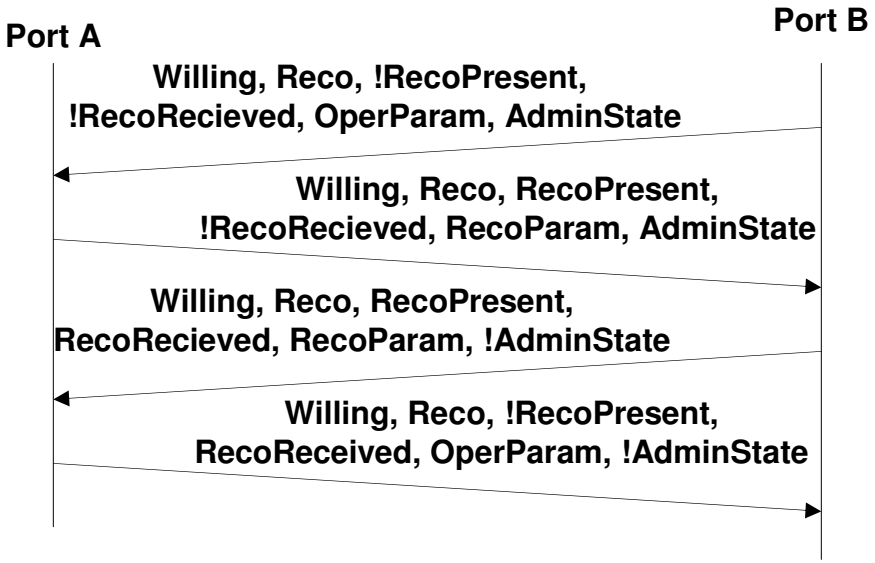

Two transfers in each direction required to transfer the parameters

Asymmetric Parameter Passing – Summary

- **Option 1 advantages:**
  - **Far fewer state variables (4 vs. 15)**
  - **Fewer state machines (1 vs. 2)**
  - **Fewer states (2 vs. 5)**
- **Option 2 advantages:**
  - **Fewer bytes in the TLVs (15 vs. 34)**

Some background on TLV space

- **Create a "reasonable worse case":**
  - **All defined TLVs enabled**
  - **Average string length of 32 octets**
  - **Management address TLV: IPv6 and 32 octet OID**
  - **Port and Protocol VLAN ID: 8 entries**
  - **Protocol Identity: 8 octets/identity, 8 entries**
  - **DCBX Application Priority TLV: 16 entries**
  - **DCBX TLV formats: per az-pelissier-dcbxtlvs-0309**
- **Assuming Option 1 for ETS: 655 octets used**
- **Assuming Option 2 for ETS: 636 octets used**

The Author's Opinion:

- **Keep it as simple as possible**
  - **But not simpler**
- **Option 1 is simpler**
  - **At the expense of 19 extra octets in TLVs**
  - **And possibly additional octets as additional TLVs that use this mechanism are defined**
- **Is Option 1 too simple?**
  - **Eventually, we may run out of space in LLDP PDUs**
    - We are in the mid 600's for one example of a "worse case"
    - Not a given we will run out of space
    - If we do, then we will need to enhance LLDP to address the space issue
      - There are a number of obvious ways to do this
  - **Will conserving 19 octets significantly affect whether we need to address this issue in the future or not?**
    - Seems quite unlikely
- **Therefore, Option 1 is not "too simple" and potentially should be favored over Option 2**

Symmetric Parameter Passing – Overview

- **Single TLV used**
  - **Always carries the current local operational state**
  - **Carries a “Willing” (“W”) bit**
  - **A port that sets the W is considered “Willing”**
  - **A Willing port shall set its operational parameter to that indicated in the received TLV if the received TLV has the W bit set to zero**

Symmetric Parameter Passing State Variables

LocalWilling: Indicates that the local port has been administratively configured to accept the parameter from the remote port. This value is included in the Willing bit of DCBX TLVs transmitted by this port.

RemoteWilling: TRUE indicates that the Willing bit was set in the last TLV received. FALSE indicates that the Willing bit was not set in the last TLV received. NULL indicates either the TLV was not present in the last LLDP PDU or that no LLDP PDUs have been received.

OperParam: The current operational state of the parameter on the local port. This value is included as the parameter in the DCBX TLV.

RemoteParam: Contains the value of the last parameter (i.e. the operational state of the remote port) received in the TLV. NULL indicates either the TLV was not present in the last LLDP PDU or that no LLDP PDUs have been received.

LocalAdminParam: The administratively configured state for the parameter. This becomes the operational state of the parameter by default, and may be overridden if the local port accepts the parameter from the remote port.

Symmetric Parameter Passing – Example TLV

TLV Type =127	TLV Info String Len=6	802.1 OUI 00-80-C2	802.1 Subtype = 11	W	Re-served	PFC Cap	PFC Enable
7 bits	9 bits	3 octets	1 octet	1 bit	4 bits	3 bits	1 Octets

Symmetric Parameter Passing State Machine

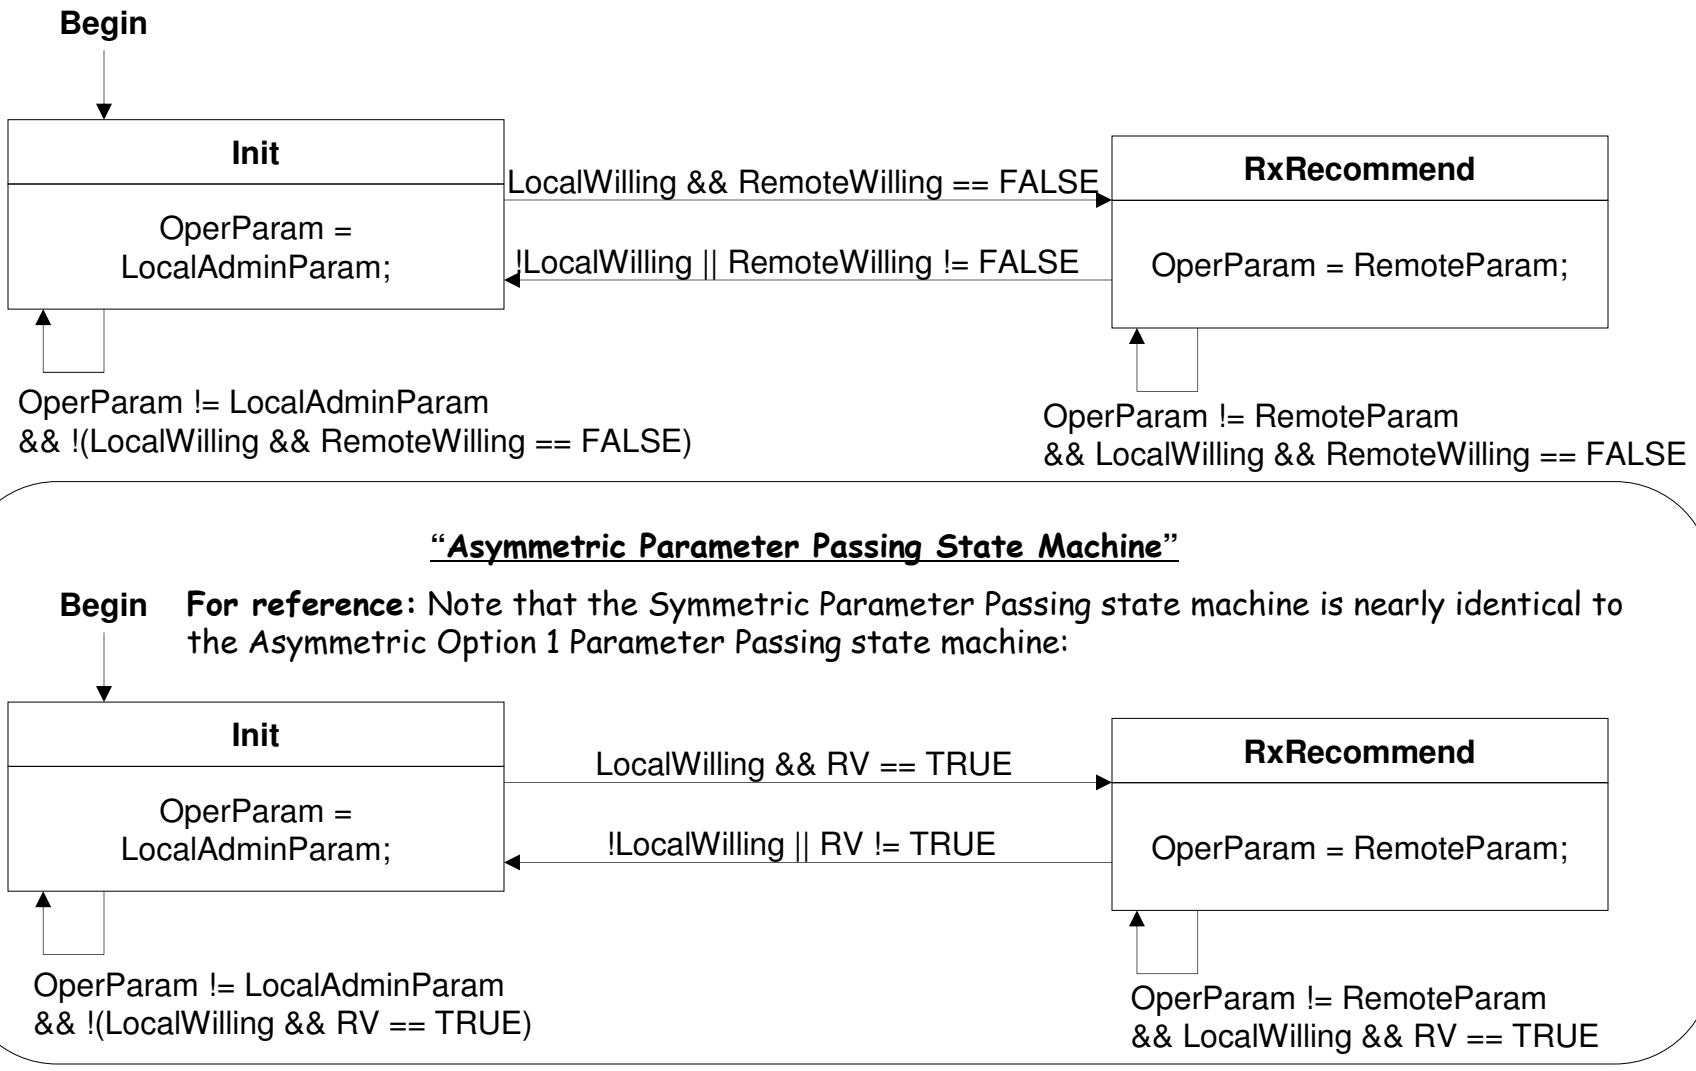

Symmetric Parameter Passing Ladder Diagram

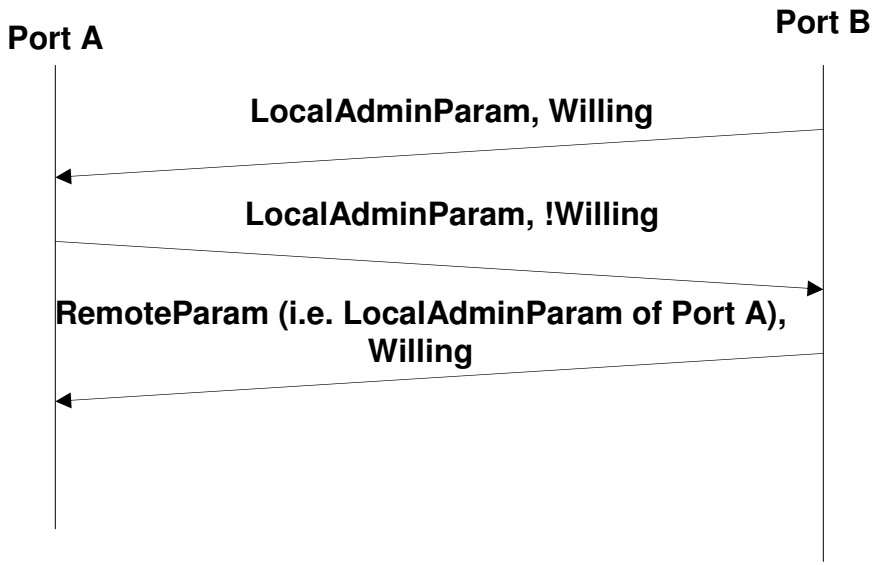

Thank You!