# White Paper:
# Control Plane Implementation on Coordinated Shared Networks (CSN)

Philippe Klein, PhD
Broadcom Corporation
philippe@broadcom.com

| | | |
|---|---|---|
| V.01 | Nov 2011 | Initial Version |
| V.02 | Jan 2012 | Add Format Figs |
| V.03 | Feb 2012 | |

## 1 Scope

While CSN data planes are well described in their respective CSN specifications, theses specifications do not describe any control plane. As long as the CSN usage was limited leaf clouds (i.e.  CSN clouds were not interconnecting other bridged LANs), this lack of specifications did not have consequence.  With the deployment of CSN “backbones” however, this lack of specifications, which leaves the control plane model to the choice of implementers, could create unsolvable interoperability issues and could impeach the support of legacy L2 network control protocols over these CSN backbones.

This white paper describes a generic scheme to implement data plane over CSN networks. This scheme allows to support legacy L2 network control protocols with minimal changes if any, to their respective protocol SW entities already developed for standard L2 Ethernet bridges.

©Copyright Broadcom 2012 – All rights reserved

## 2 CSN Network Definition

A CSN is a contention-free, time-division multiplexed-access network, supporting reserved bandwidth based on priority or flow (QoS). One of the nodes of the CSN acts as the Network Coordinator (NC) node, granting transmission opportunities to the other nodes of the network. The NC node also acts as the bandwidth resource manager of the network.

Current deployed CSNs are built on two physical technologies : coaxial and powerlines , respectively specified in the following standard MoCA™ and IEEE 1901. Coordinated Shared Network (CSN) characteristics

CSNs support two types of transmissions: unicast transmission for node-to-node transmission and multicast/ broadcast transmission for one-node-to-other/all-nodes transmission. Each node-to-node link has its own bandwidth characteristics which could change over time due to the periodic ranging of the link. The multicast/broadcast transmission characteristics are the lowest common characteristics of multiple/all the links of the network.

A CSN network is physically a shared network, in that a CSN node has a single physical port connected to the half-duplex medium, but is also a logically fully-connected one-hop mesh network, in that every node could transmit to every other node using its own profile over the shared medium.

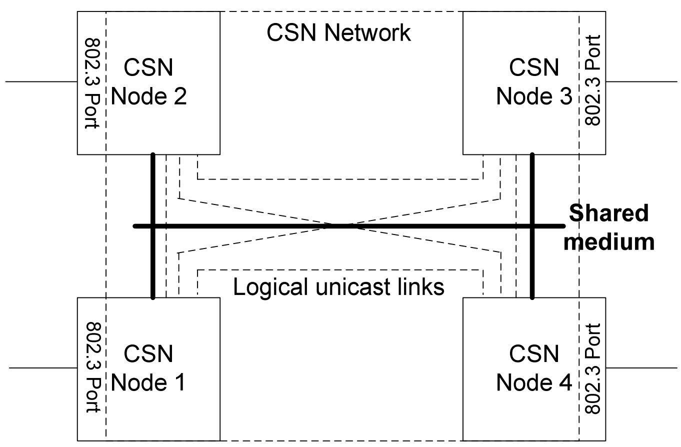


**Figure 1: CSN Network**

## 3 Backbone CSN

CSN Networks transport encapsulated standard 802.3 frames over their networks. CSN nodes export a 802.3 port at the edge of the CSN network. From the L2 IEEE 802.1 stand point, a CSN network is modeled as a distributed Bridge.

©Copyright Broadcom 2012 – All rights reserved

# 4 CSN Data Plane

A CSN's Data Plane behaves identically to a regular L2 Bridge for Unicast, Multicast and Broadcast data frames. In a CSN network, each node is equivalent to a Bridge's port. Each node-to-node link is equivalent to a Bridge's path from an ingress port to an egress port.

While in a standard bridge, the forwarding engine and the forwarding database are single entities, in a CSN the forwarding engine and the forwarding database are distributed in every node.

CSN's specifications define how CSN nodes maintain their own bridging table to map the MAC DA address of 802.3 frames bridged by the network to the CSN addresses of the egress nodes bringing theses addresses.

CSNs forward frames with unknown mapping to all the egress ports but some CSN specification could require shaping these propagations with a limiting frame rate to avoid broadcast storms.

# 5 CSN Control Plane

Contrary to the data plane, control plane are not described in the existing CSN specifications and network control data units are currently handled as regular data frames making the CSN transparent to these protocols.

Although seen as a bridge from the data path stand point , the CSN will not support legacy loop protection protocol such Spanning Tree Protocol or bandwidth reservation protocol such RSP.

## 5.1 Central vs. Distributed Control Plane

A first option to implement a control plane in a CSN network would be to include a control plane entity in each CSN node, in which case each ingress node would act as an independent bridge.

This scheme presents several drawbacks:

1) If the CSN network includes a large number of nodes, the control traffic overhead between the ctrl planes entities will be significant.

2) For a given link, the bridging database from each end will not be shared (unless a new protocol is developed between both ctrl plane entities to share information).

The second option, which duplicates the scheme of legacy bridge, implements a single control plane entity within the CSN network.

©Copyright Broadcom 2012 – All rights reserved

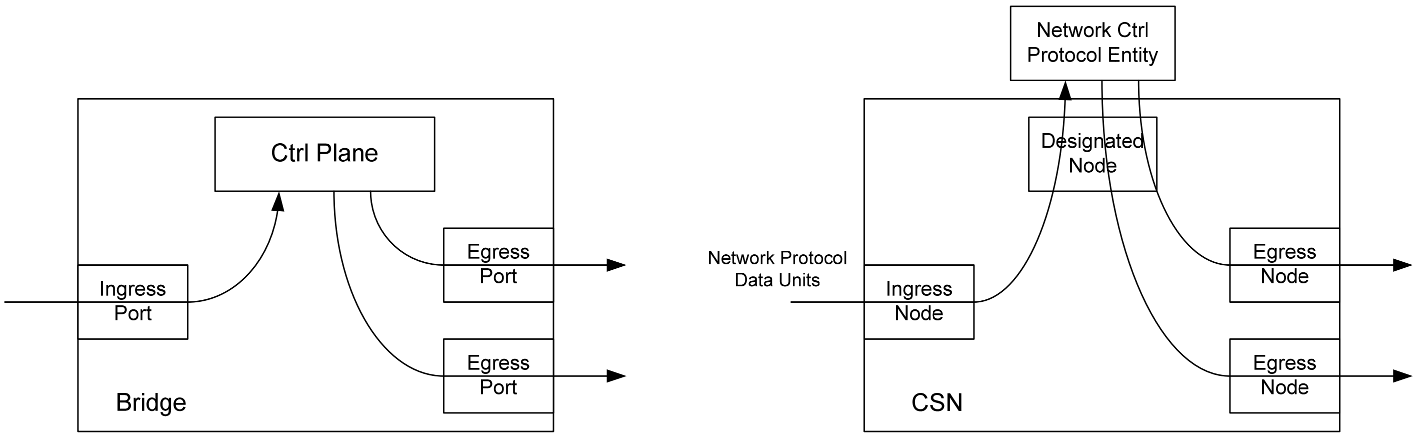


**Figure 2: Single Control Plane emulation on CSN distributed bridge**

## 5.2 Designated Node

Let's call the node on which the single plane entity runs, the Designated Node (DN). A CSN network could either include a single Designated Node for all the supported L2 network protocols or could select a node on a **per protocol or service base.**

### 5.2.1 DN Selection and Migration

Depending on the CSN technology, the DN might correspond to a static node or dynamically migrate between nodes during normal operation.

### 5.2.2 Flexibility

The DN scheme allows flexibility as a given node may be the DN for many protocols/services/features, but the node is not necessary required to implement the control and functions for all these applications.

#### *5.2.2.1 DN Selection Simple Algorithm*

Although the DN selection could be network specific, this paragraph describes a single DN selection algorithm for a given network control protocol/service called "NCP"):

1) A node broadcasts a DN selection request to all the nodes to query
   - which nodes are "NCP_Aware" (i.e. nodes able to deliver NCP Data Units (NCPDUs) to the NCP_Designated_Node and handling NCP related commands sent by the DN if any)
   - which nodes are "NCP_DN_Capable" (i.e. nodes able to deliver NCPDUs to a NCP entity)
   - and which node is the "NCP_Designated_Node (NCP_DN)"

   in two cases:
   a) the node is a new node to the CSN network
   b) the previously selected NCP_DN has been removed from the CSN network topology list (an information present in CSN networks)

2) The NCP_DN_Capable nodes broadcast their responses to the DN selection request:

©Copyright Broadcom 2012 – All rights reserved

a) if one NCP_DN_Capable node indicates it is also the NCP_DN , no further action is required.

b) if none of the NCP_DN_Capable node indicates it is the NCP_DN, the NCP_DN_Capable node with the higher CSN specific node ID becomes the NCP_DN (i.e. the node initializes the NCP Service daemon) and indicates it to the other NCP _DN_Capable (and NCP_Aware) nodes (although there is no ambiguity nor race condition between NCP_DN_Capable nodes due to the uniqueness of their node IDs).

### 5.2.3 DN Database

Similary to regular bridge, the NCP_DN over time constructs its NCP database by handling the "NCPDUs" but this database is not a permanent database. If the NCP_DN migrates, the new NCP_DN could dynamically reconstruct the NCP database.

Therefore a NCP_DN failover when a NCP_DN fails could be handled without any need for duplication and backup database mechanism between the NCP_DN_Capable nodes.

### 5.2.4 Fault Tolerance

The DN model is fault tolerant:

- DNs could be re-selected dynamically,
- Bridge protocol's database are not persistent and could be dynamically reconstruct by the new DN,

Fault tolerance could be increased by using a multi-DN configuration in which different DNs could be selected per service / per protocol basis

### 5.2.5 "Mixed" Configuration

The DN scheme could be deployed on a CSN with both NCP_Aware/DN_Capable and legacy nodes. The NCP_DN Control Plane will only control the NCP Aware/DN_Capable nodes. The other "unmanaged" nodes will be seen as external bridges as described in Figure 3.

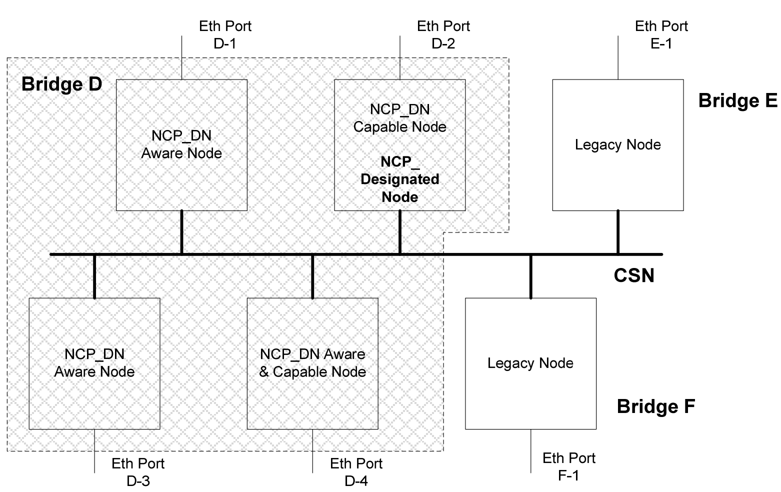


©Copyright Broadcom 2012 – All rights reserved

**Figure 3 : DN Scheme in a mixed CSN environment**

# 6 NCPDU handling on a CSN

A NCP-aware CSN node identifies NCPDUs received on its non-CSN interface (either the interface to another network media or to upper layers of the node) by their destination MAC address (**Nearest Customer Bridge Group Address**) and NCP EtherType.

1) Non-DN nodes send NCPDUs to the DN over the CSN.

2) The DN delivers NCPDUs, along **with information about the originating interface**, to the NCP Service.

3) The NCP Service on the DN node handles the NCPDUs which could trigger invocations of CSN's Specific primitives (see as examples, the MSRP and RSTP examples in paragraph 7 & 8)

**4)** The DN node could in turn sends NCPDUs to the CSN egress nodes along **with information about the egress interface.**

## 6.1 DN Adaptation Layer

**The DN Adaption Layer shall as illustrated in Figure 4**

Figure 4 :

1) unbind the NCPDU and the network topological information (ingress port) from the ingress MSDU received from the DN's egress port and present these parameter to the NCP Service entity standard interface

2) bind the NCPDU and the network topological information (egress port) received from the NCP Service entity standard interface into a single MSDU presented to the DN ingress port.

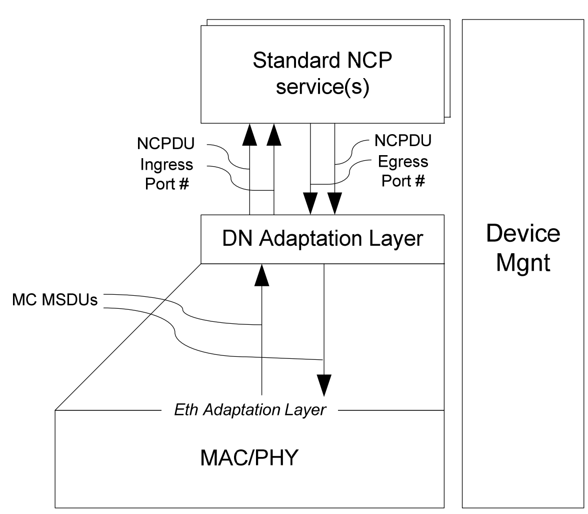


**Figure 4: DN Adaptation Layer**

# 7 802.1Q MSRP Example

The single control plane scheme described in this document has been selected for handling the MSRP protocol over CSN and is specified in the IEEE Std 802.1Q-2011 Specifications, Annex C) and illustrated by Figure 5-A and Figure 5-B:

*a) A MSRP-aware CSN node identifies MSRPDUs received on its non-CSN interface (either the interface to another network media or to upper layers of the node) by their destination MAC address and EtherType.*
*b) Non-DMN nodes send MSRPDUs to the DMN over the CSN.*
*c) The DMN delivers MSRPDUs, along with information about the originating interface, to the MSRP Service.*
*d) The DMN translates the MSRP TSpec parameters into CSN QoS parameters and invokes the CSN's Protocol Specific QoS transactions with the CSN Network Coordinator as follows:*
   *1) When the DMN receives a Talker Advertise message originated from an upstream CSN node, the DMN invokes a bandwidth query transaction with the CSN Network Coordinator to check whether or not the bandwidth advertised in the message's TSpec is available on each upstream to downstream node link of the CSN network. In addition the DMN associates the MSRP attributes with the CSN's native QoS records via the StreamID (which is included in the CSN frames).*
   *2) When the DMN receives a Listener Ready message originated from a downstream CSN node, the DMN invokes a bandwidth reservation transaction with the CSN QoS manager to reserve the bandwidth associated with the message's StreamID on the downstream to upstream CSN node link.*
*e) After the DMN completes the CSN QoS transactions, the DMN behaves as an MSRP application on a Bridge and propagates (MAP) and distributes (MAD) MSRP attributes.*

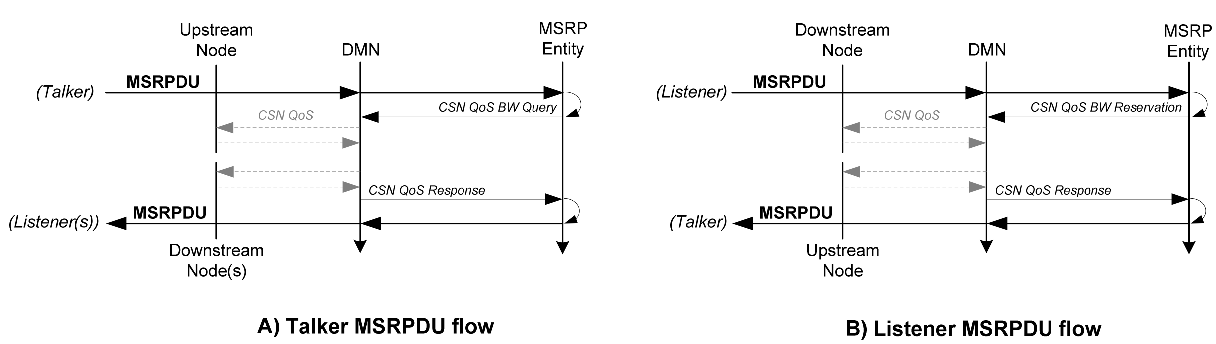


**Figure 5: MSRPDU Flows within the CSN**

Figure 6 illustrates the MSRP implementation over a distributed CSN bridge.

©Copyright Broadcom 2012 – All rights reserved

*Note: The comparison between Figure 6 and Figure 8 shows the similarity of the handling applied to multiple protocols.*

# 8 RSTP Example

As a second example, the following section describes how the standard RSTP protocol could be implemented on a CSN to protect the CSN network against loops.

( As mentioned previously this control plane scheme is not limited to a specific protocol but could easily be extended to any L2 network control protocols)

*Note: (x) in the flowing paragraph refers to the same index in* Figure 6

RSTP messages (BPDU) are identified by their MAC Destination Address & Ethertype:

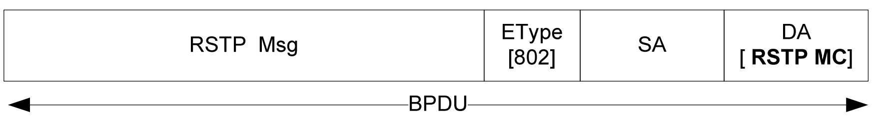


Upon receiving a BPDU **(1)**, a CSN ingress node discriminates this BPDU based on its MC DA address and Ethertype and handles it thru its “Ctrl Path”. The ingress node:

1) appends its own CSN node_ID as Port Number to the BPDU: **(2)**

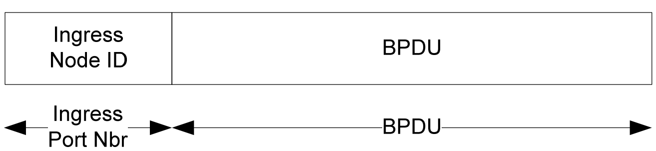


2) encapsulates the resulting [node_ID | BPDU] frame into a CSN **unicast** MSDU addressed to the DRN (Designated RSTP Node) **(2)**

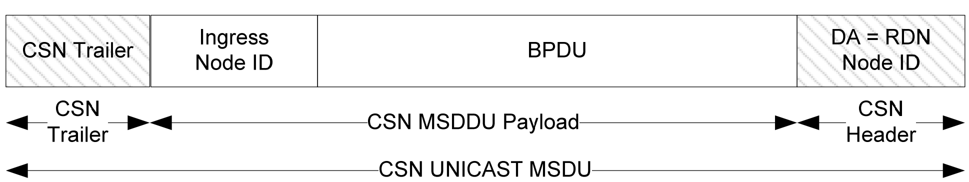


3) transmits this MSDU **(3)**

4) Upon reception over its CNS interface, the **DRN** handles the MSDU as any regular data frame:

   The MSDU Payload is de-capsulated from the CSN network’s MSDU and transmitted through the node’s Ethernet layer interface.

©Copyright Broadcom 2012 – All rights reserved

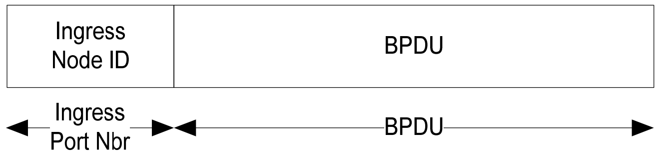


The frame is eventually delivered by the protocol stack to the RSTP Service daemon **(4)** (the RSTP Service daemon is registered to the specific multicast address / Ethertype of the BPDU).

The ***DRN adaption layer*** of the RSTP Service daemon extracts the standard BPDU and the topological info (i.e. the ingress port this BPDU was received on) and presents them to the **standard RSTP protocol entity**.

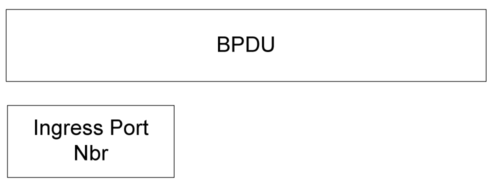


The RSTP Service deamon interacts with the underlying CSN network in two ways:

**a) propagates BPDU over specific egress port(s) :**

To propagate BPDU over specific egress ports, the DRN adaption layer of the RSTP Service daemon, appends the egress node_ID to the BPDU **(7)** and transmit the resulting frame to the DRN **(8)**.

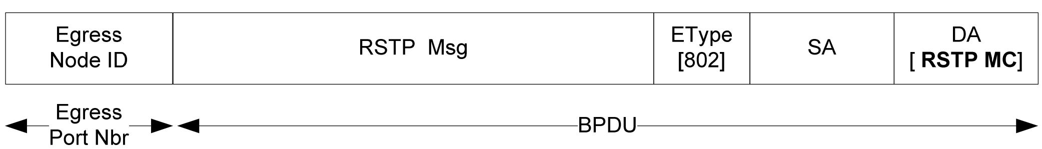


Upon receiving the frame from the RSTP Service deamon over its Ethernet interface, a DRN discriminates this BPDU based on its MC DA address and Ethertype and handle it thru its “Ctrl Path”. The DRN:

1) retrieves the BPDU **(9)**

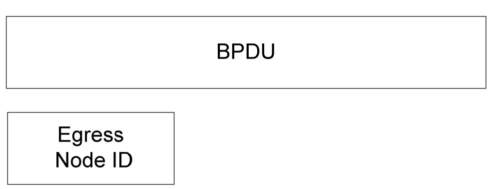


**2)** encapsulates the BPDU into a CSN **unicast** MSDU addressed to the egress node_ID transmitted by the RSTP Service deamon **(9)**

©Copyright Broadcom 2012 – All rights reserved

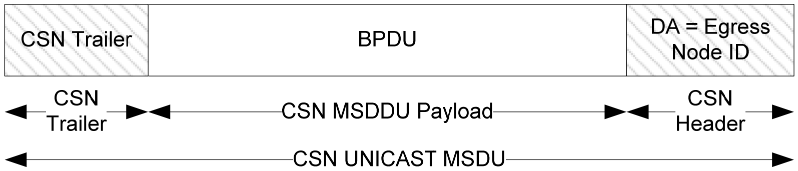


3) transmits this MSDU **(10)**

The receiving egress node handles the MSDU as any regular data frame:

The [ BPDU] is de-capsulated from the network's MSDU and transmitted through the node's Ethernet layer interface **(11)**.

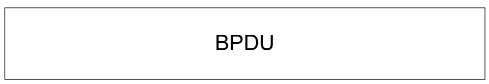


**b) Blocked Port state(s) to prevent further forwarded traffic**

The RSTP Service daemon sends "blocked_port" RSTP port state command [format TBD {node ID bitmask, blocked_bitmask}] to the nodes of the CSN network through out-of-band commands.

The RSTP Service daemon invokes an I/O Ctrl function **(7')** (could be a MLME primitive or any implementation specific command) to request the DRN to transmit a block_port command to the other (non-DRN) CSN nodes through the CSN network specific ctrl message protocol. **(10)**

The blocking mechanism is implementation specific to each node.

©Copyright Broadcom 2012 – All rights reserved

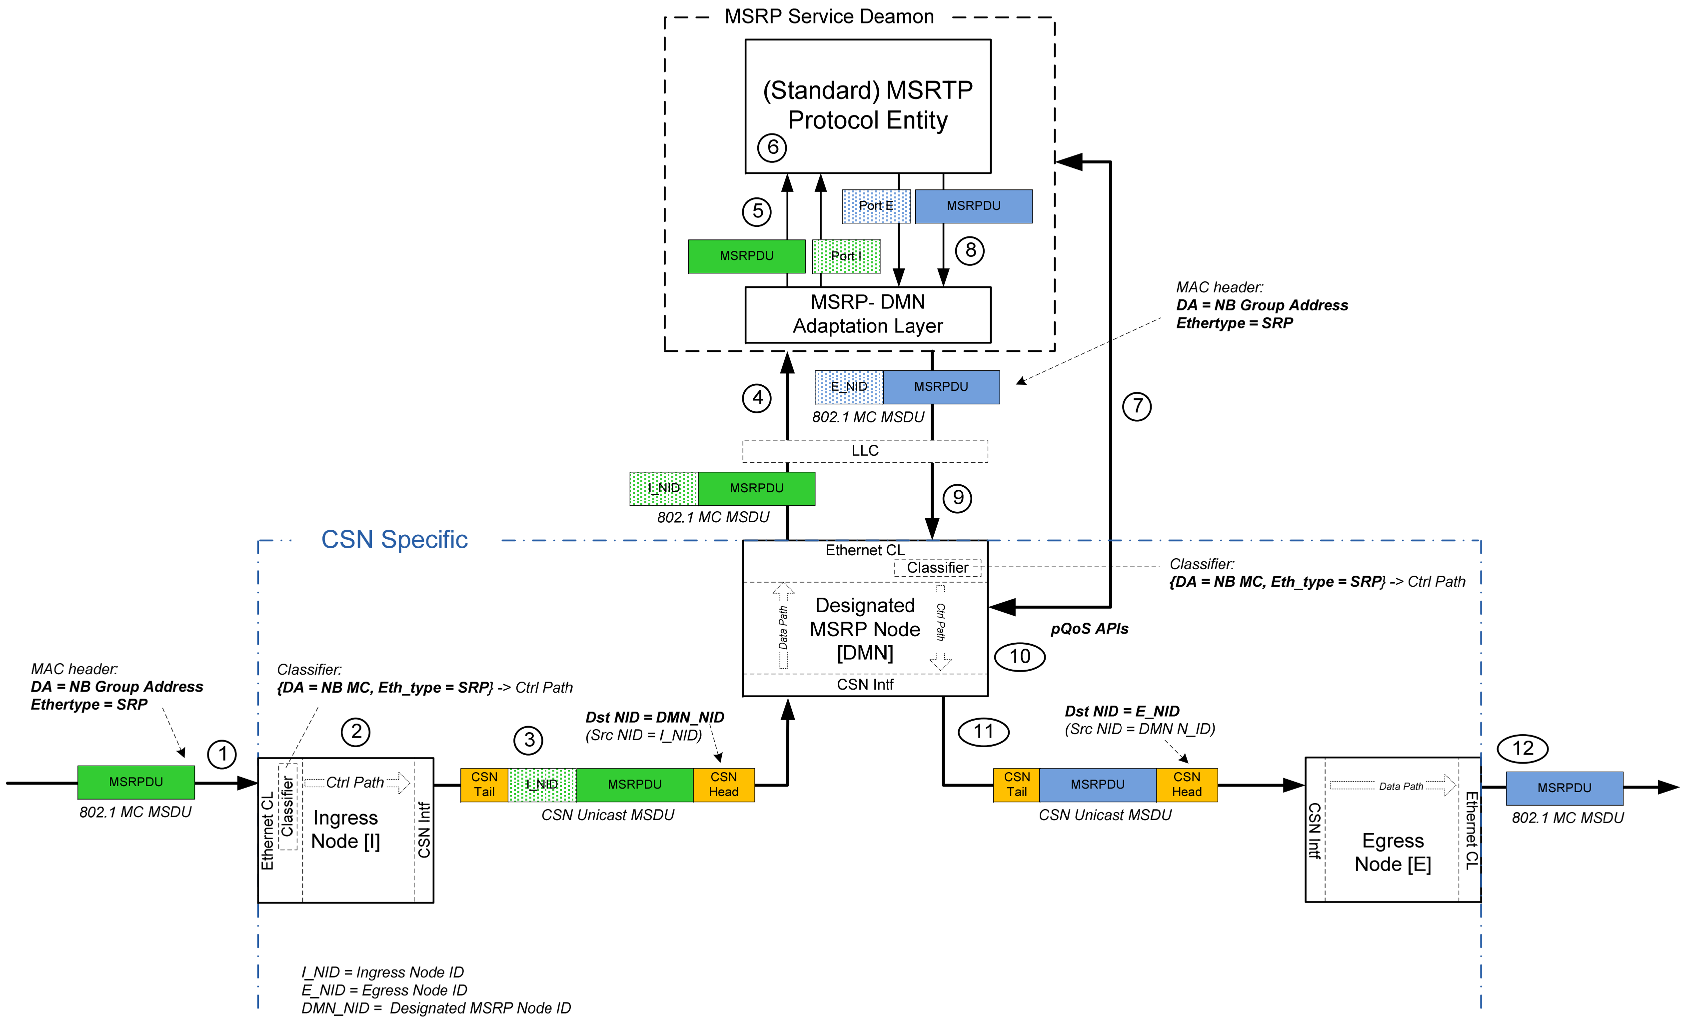


**Figure 6: MSRP implementation over a distributed CSN bridge**

©Copyright Broadcom 2012 – All rights reserved

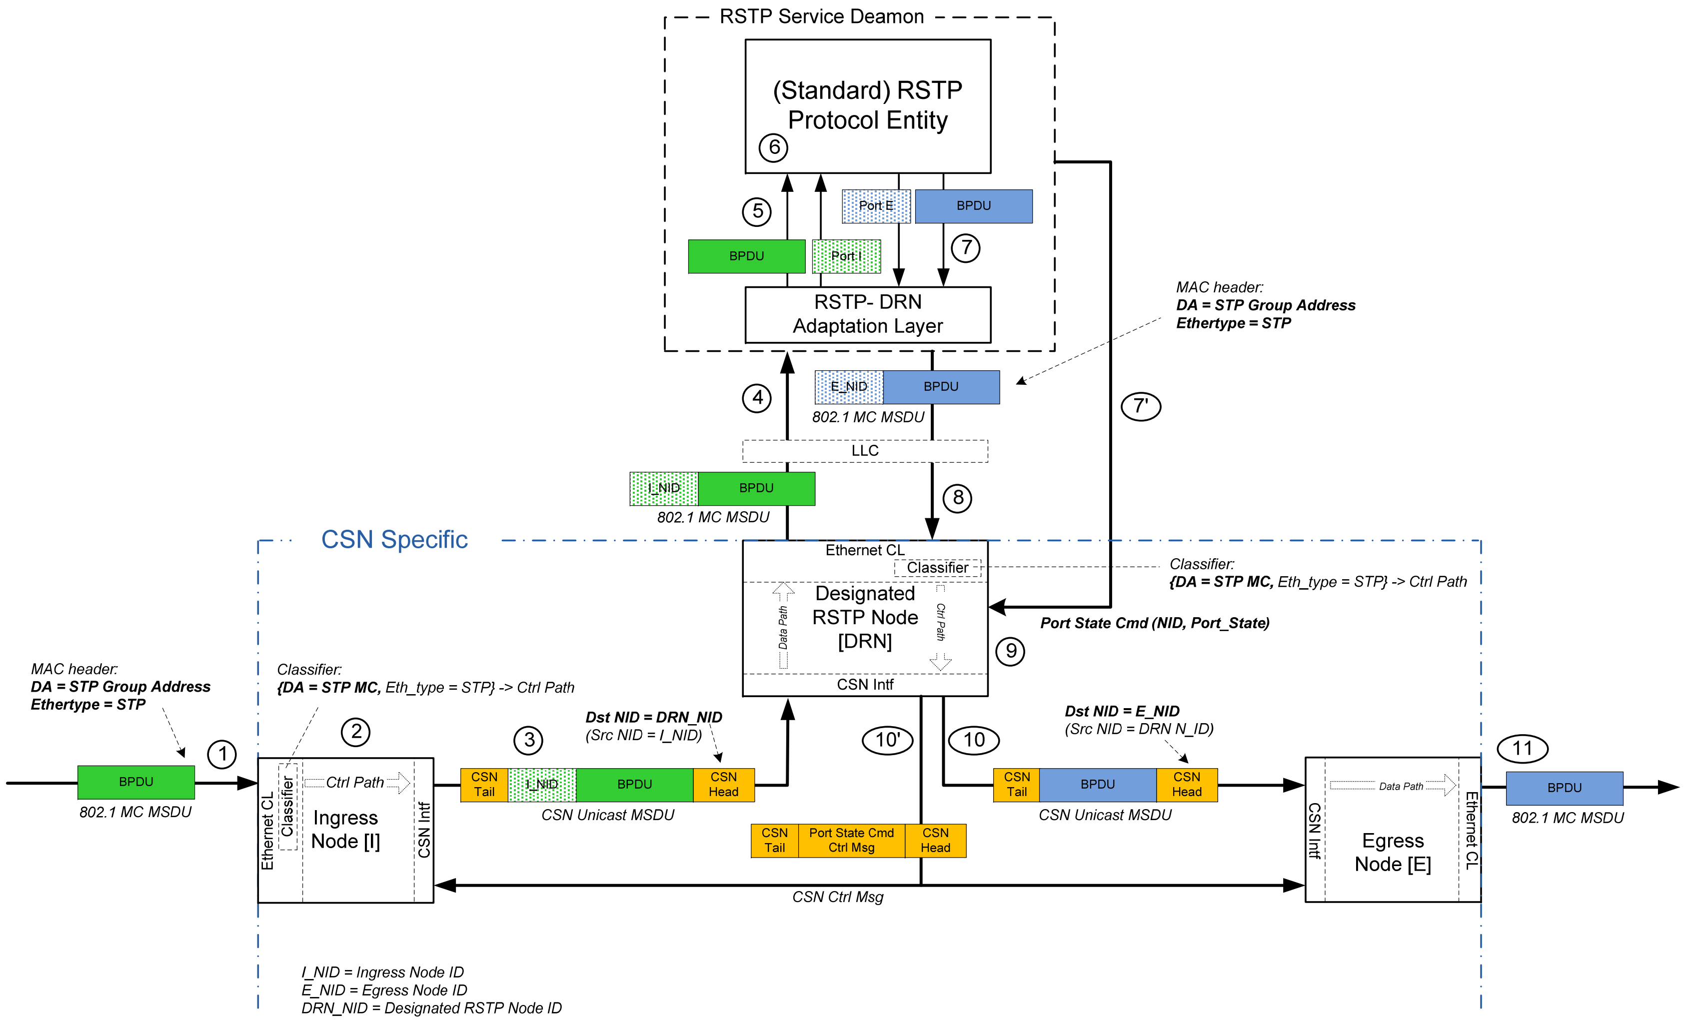


**Figure 7: RSTP implementation over a distributed CSN bridge**

©Copyright Broadcom 2012 – All rights reserved