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

Philippe Klein, PhD
Broadcom Corporation
V0.1 - Nov 2011

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 were limited to be leaf cloud (i.e. CSN clouds were not interconnecting Control plane are left to the choice of other bridged LANs) this lack of specifications did not had consequence. With the deployment of CSN "backbone" however, this lack of specifications which leaves the control plane model to the choice of implementers, could create unsolvable create interoperability issues and could impeach the support of legacy L2 network control protocols over the 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 bridge.

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.

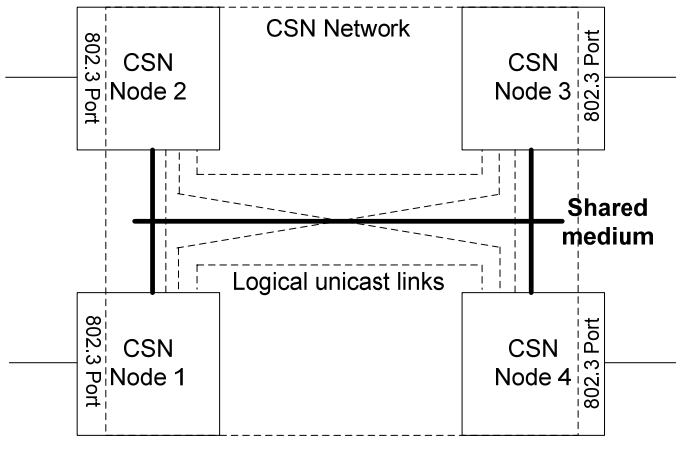

Figure 1: CSN Network

3 Backbone CSN

CSN Networks transports 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. CSN's Data Plane

4 CSN Data Plane

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.

CSN Control Plane supports the exact same functionally that regular Bridge.

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 require to shape these propagations to a low 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 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 present several drawback:

1) The first one is that the CSN network includes a large number of nodes , the control traffic between the ctrl planes entities will be significant.

2) The second drawback if that for a given link , the bridge database on each ends 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, is to implement a single control plane entity within the CSN network.

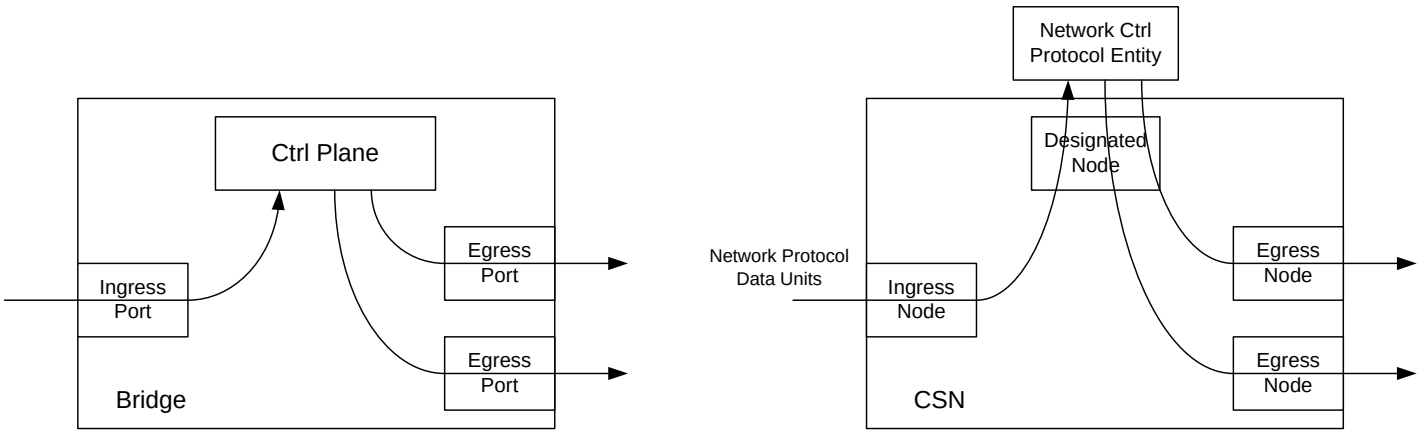

Figure 2: Single Control Plane emulation on CSN distributed bridge

5.2 Designated Node

Let's called 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 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.

5.2.1.1 DCN Selection simple algorithm

Although the DN selection could be network specific, this paragraph describes a single DN selection algorithm (let's called a given network control protocol NCP):

1) A node broadcasts a DN selection request to all the nodes to query

- which nodes are “DN_Capable” (i.e. able to deliver NCPDUs (NCP Data Units) to a NCP entity)
- and which node is the Designated Node

in two cases:

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

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

a) if one DN_Capable node indicates it is also the DN Node, no further action is required.

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

5.2.2 DCN Database

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

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

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.

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

3) 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 an example IEEE Std 802.1Q-2011 specifications, Annex C or the RSTP example in paragraph 8)

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

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)

8 RSTP Example

As an 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 3

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 node_ID to the BPDU: [node_ID | BPDU] **(2)**

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

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

4) Upon reception over its CNS interface, the **RDN** (RSTP Designated Node) handles the MSDU as any regular data frame:

   The [node_ID | BPDU] is de-capsulated from the network's MSDU and transmitted through the node's Ethernet layer interface. The [node_ID | BPDU] 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 ***RDN 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 RDN adaption layer of the RSTP Service daemon, appends the egress node_ID to the BPDU **(7)** and transmit the resulting frame to the RDN **(8)**.

Upon receiving a [node_ID | BPDU] frame from the RSTP Service deamon over its Ethernet interface, a RDN discriminates this BPDU based on its MC DA address and Ethertype and handle it thru its "Ctrl Path". The RDN:

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

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

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 RDN to transmit a block_port command to the other (non-RDN) CSN nodes through the CSN network specific ctrl message protocol. **(10)**

The blocking mechanism is implementation specific to each node.

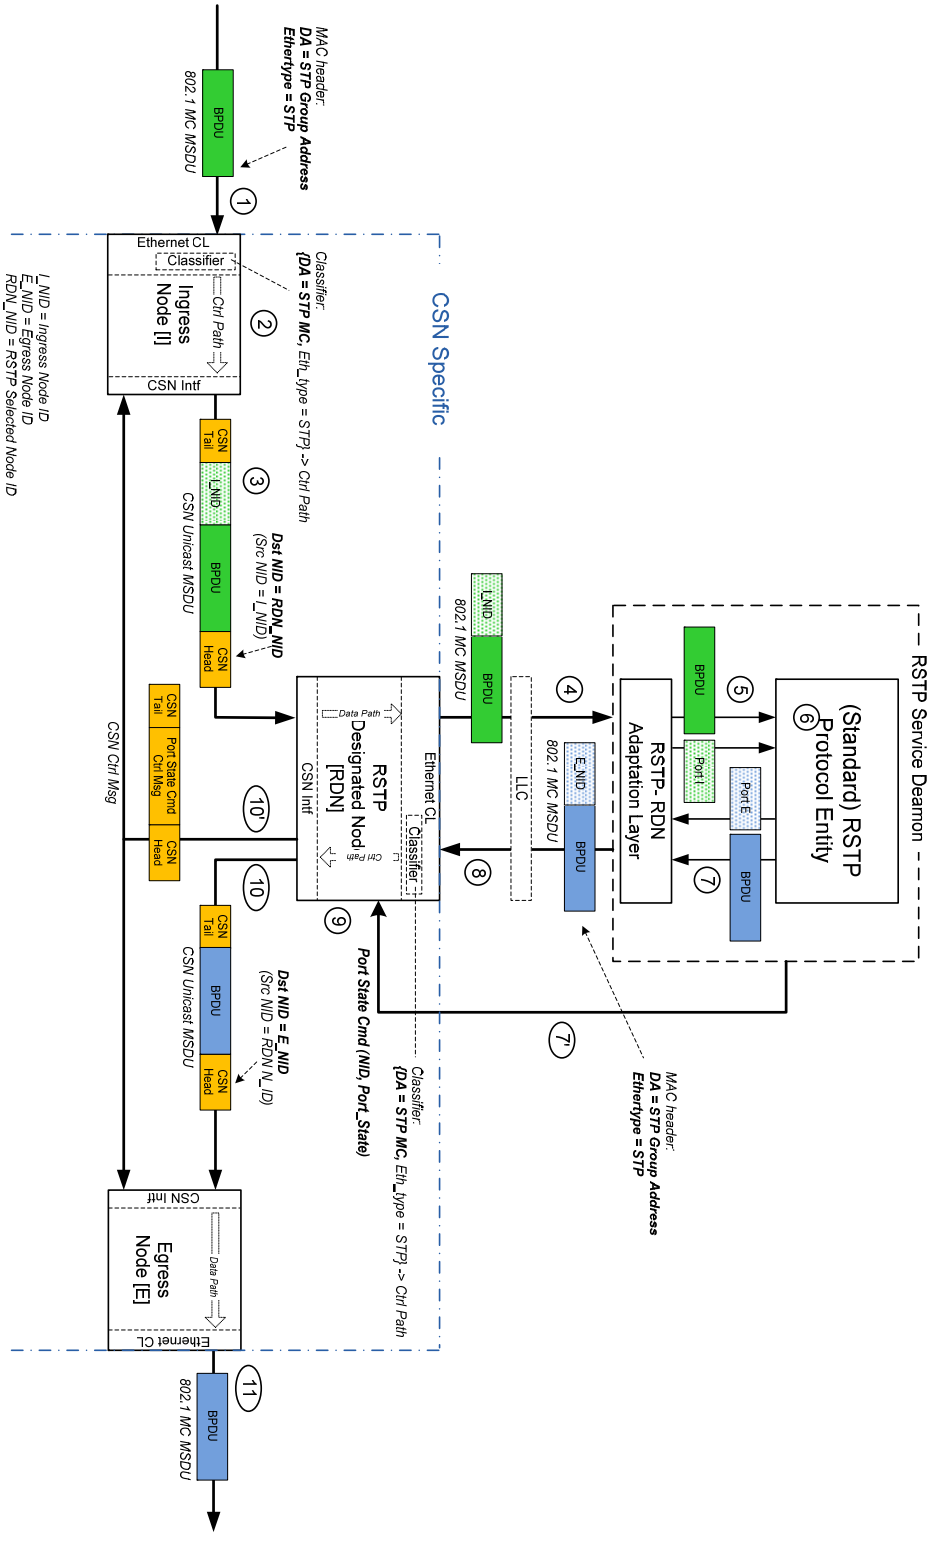

Figure 3: RSTP implementation over distributed CSN bridge