802.1AS Hot Standby Amendment: Scope Discussion

Rodney Cummings
National Instruments

Introduction

- November proposal to move 802.1AS hot standby from P60802 to new amendment of 802.1AS
  - Motion approved to work on PAR in January meeting
- Important to limit scope of amendment PAR
  - Avoid delay to dependent P60802 project
  - Hot standby topic is prone to proposals that seem simple at first, but later turn out to be complex
- Question for group
  - Are we willing to scope to what P60802 needs, or do we want a complete "reset" on the hot standby topic?
    - (this presentation assumes scope to P60802)

Include in Scope

- Specifications in P60802/D1.1 for each time (i.e. Working Clock or Global Time)
  - 2 gPTP domains only (Table 2)
  - Function to "merge" 2 domains into 1 time
  - Support for ARB timescale (not globally traceable)
  - BMCA disabled (externalPortConfiguration=true)
- Presumed consensus for P60802
  - Domain quality (also known as "synced")
    - Line 1079 of P60802/D1.1: Determines whether domain can be used
    - Mechanisms specified in 802.1AS amendment (e.g. offsetFromMaster)
    - Numbers for each mechanism specified in P60802
      - (e.g. offsetFromMaster shall be less than 100 ns)

Include in Scope (Presenter's Opinion)

- Function to "split" 1 domain into 2 domains
  - Shown in 2014 presentation
  - Sometimes referred to as "coupled rings" use case
  - One network with multiple "segments" (e.g. rings)
  - Disjoint segments are separated by 2 relays
  - Requirements
    - Support 1 GM failure per network (all segments)
    - Support 1 link or relay failure per segment
      - I.e. for three segments, three links can fail
  - Assumes that over a large network, failure of a link/relay is more likely than failure of a GM
    - Implies more focus on hardware failures than software failures

Exclude from Scope (Presenter's Opinion)

- Mitigation of byzantine fault
  - Requires at least 3 domains, voting algorithm, and possibly hardware requirements
- BMCA
  - Good: Replace faulted GM, merge >1 physical network, etc
    - E.g. Figures 7-1 and 7-2 of 802.1AS-2020
  - Bad: With transient faults, BMCA "flaps" between GMs/trees
    - Perceived to be unpredictable; Difficult to achieve stable sync
  - Make a better BMCA?
    - Starts simple, but typically turns out to be complex

Concern: Examples in 802.1AS-2020

- Figures 7-4, 7-5, and 7-6 show examples of hot-standby
- Examples were designed to show what is possible
  - They serve that purpose well
- Don't necessarily represent practical implementations

Figure 7-4 and 7-5

- Align with P60802/D1.1, but subset of full support
- 7-4 has 2 GMs in one box
  - If box fails, time fails
- 7-5 has 1 sync tree (no ring)
  - If link fails, time fails

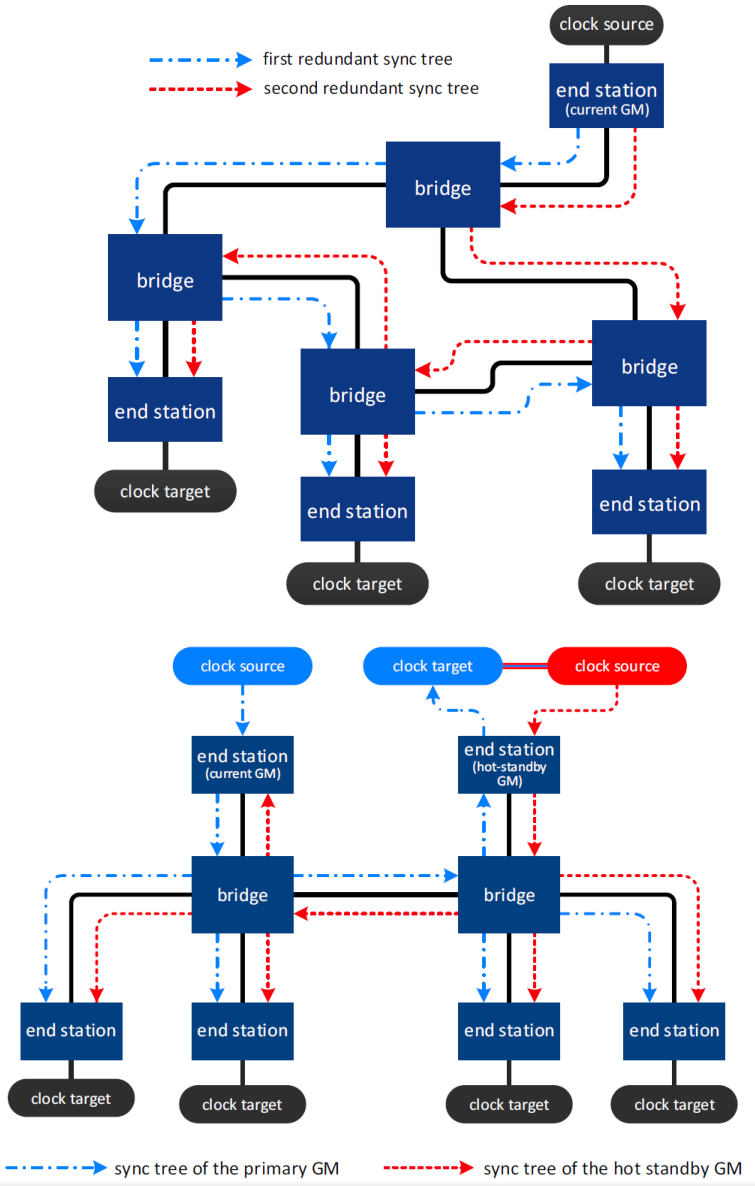

Figure 7-6

- Full support, but uses 4 domains
- Trees redundant, but that is possible with 2 domains

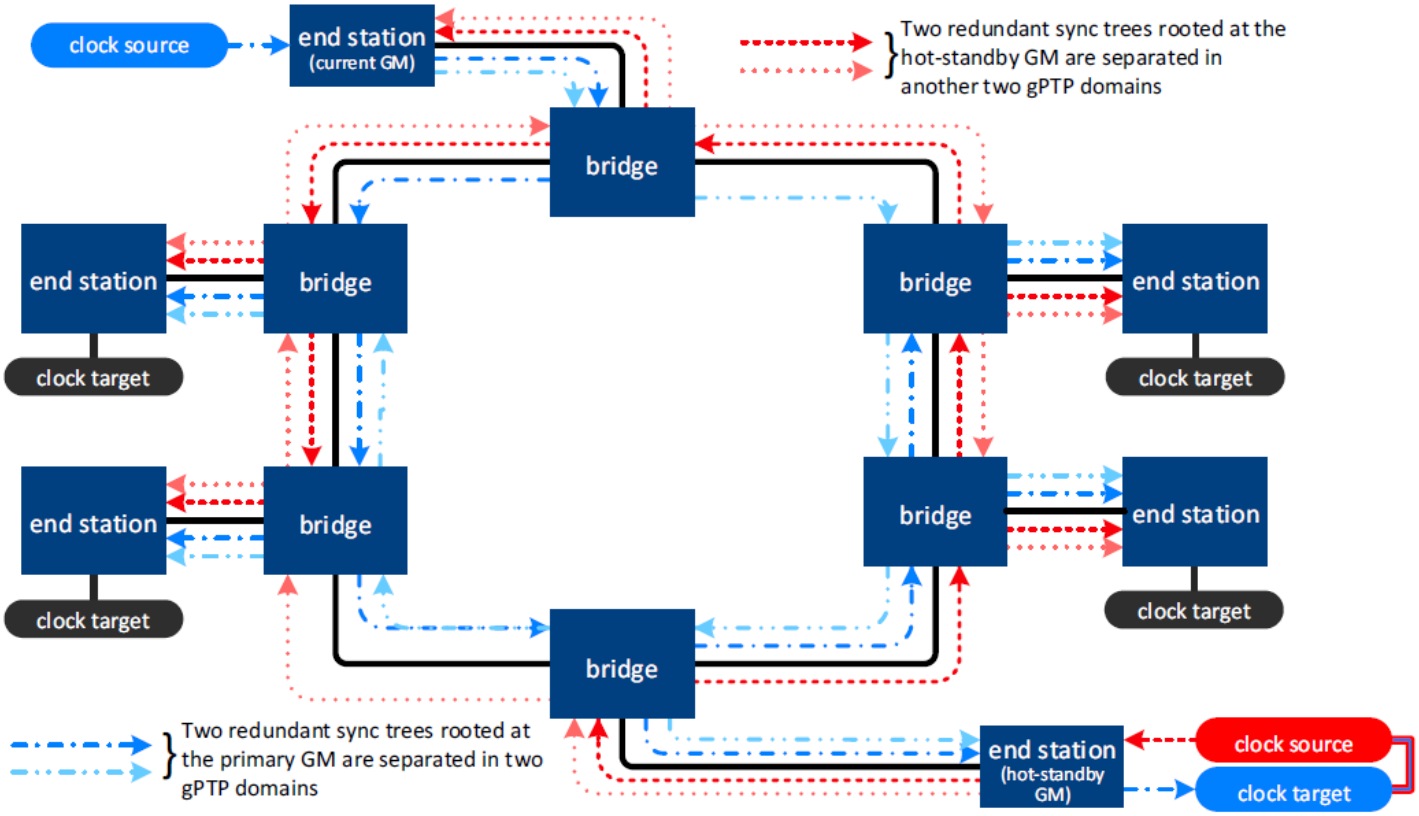

Concerns for Figure 7-6

- Pro: Amendment should cover all 3 examples
  - 802.1 put work into these examples
- Con: Figure 7-6 not possible with P60802/D1.1 Table 2
- Con: Missing practicality could misdirect project
  - 4 domains are likely to "reset" discussion of topic
- Con: 4 domains twice as expensive as 2 domains
  - Assuming expense includes software development, testing, diagnostics in field, etc

Proposed Approach for Scope

- Most IEEE PAR scopes describe what is included
  - Since it is a scope,
    everything not explicitly included,
    is excluded
  - We are accustomed to this technique in 802.1

- If scope states that externalPortConfiguration=true, that excludes BMCA

How to Exclude Byzantine Fault?

- Two possible approaches
  1. Scope states 2 domains only
     - Pro: Since byzantine mitigation needs 3 domains, it is excluded
     - Con: Excludes other uses of domains like Figure 7-6
  2. Explicitly state that byzantine fault mitigation is excluded
     - Pro: Explicit and clear
     - Con: Not the typical PAR convention
- Proposed scope uses the 1st approach

Scope Proposal

"This amendment specifies procedures and managed objects for hot standby redundancy, including:

- Function to merge two domains into one time.
- Function to split one domain into two domains.
- Specification of mechanisms that determine whether a domain has sufficient quality to be used for hot standby.
- The externalPortConfiguration variable is true for all hot standby domains.
- Hot standby domains support the arbitrary timescale.
- Change existing text for hot standby to align with new features.

"