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

Include in Scope

- Specifications in P60802/D1.1
  - 2 domains only
  - Function to "merge" 2 domains into 1 time
  - Support for ARBitrary timescale
  - 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
    - Numbers for each mechanism specified in profile (i.e. P60802)
  - Automatic restoration of a domain can be done externally
    - E.g. by CNC, NMS, SDN controller

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

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

Proposal: Scope to Examples

- Proposal is to limit scope to examples only, such as "Fully specify redundancy as shown in the examples of Figures 7-4, 7-5, and 7-6 in IEEE Std 802.1AS-2020."
- Next slides discuss these examples...

Figure 7-4

- Shows
  1 GM,
  2 trees,
  using 2 domains
- Interesting as an example of what is possible
- Not relevant in practice
  - Using 2 domains, there is no benefit to avoiding 2 GMs
  - P60802/D1.1 assumes redundant GMs and trees

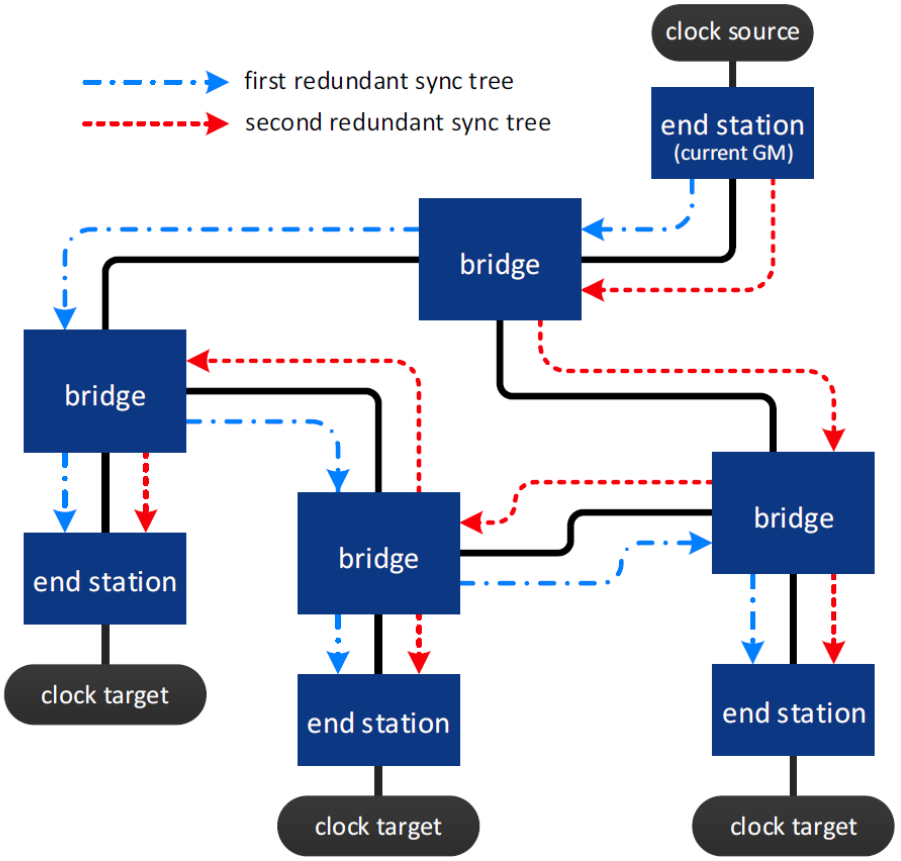

Figure 7-5

- Shows
  2 GMs,
  2 trees,
  using 2 domains
- Trees are not redundant
  - No ring/mesh
- Arguably not practical
  - Models P60802 specs, but missing important tree redundancy
  - Example incrementally adds another possibility

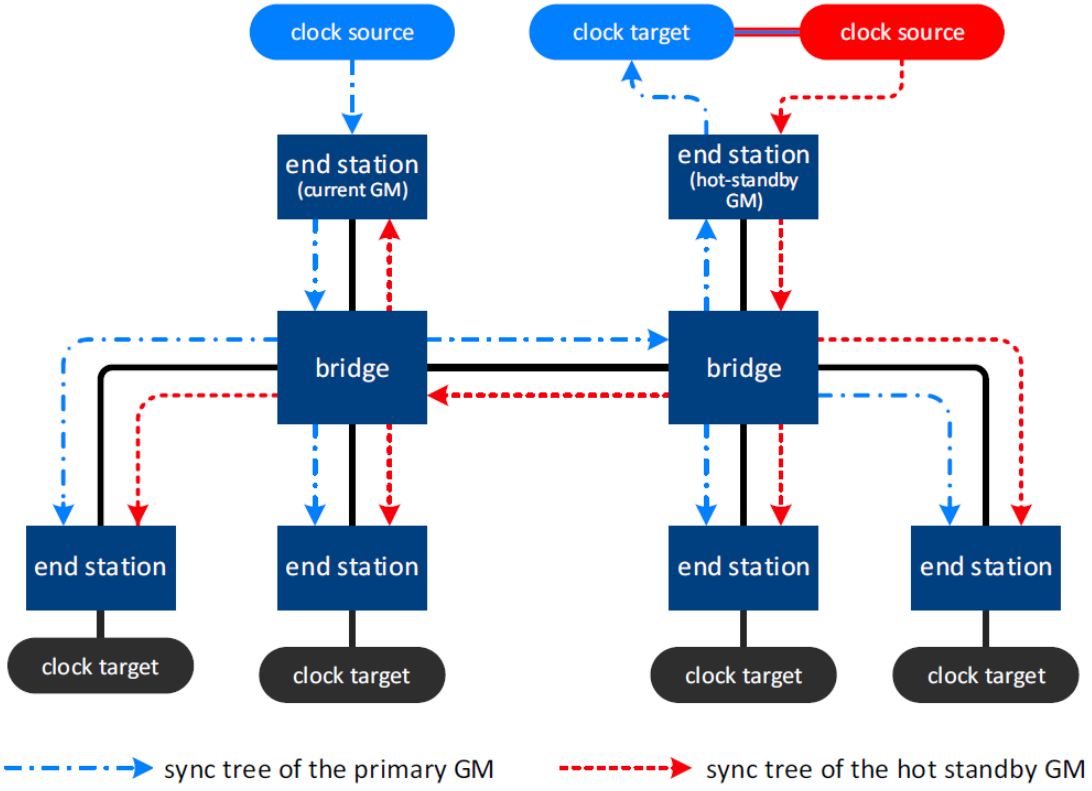

Figure 7-6

- Shows
  2 GMs,
  4 trees,
  using 4 domains
- Trees redundant, but that is possible with 2 domains
  - 4 domains can improve performance by mitigating rate ratios
- Arguably not practical
  - Not possible with P60802/D1.1 (only 2 domainNumber values)
  - Example incrementally adds another possibility

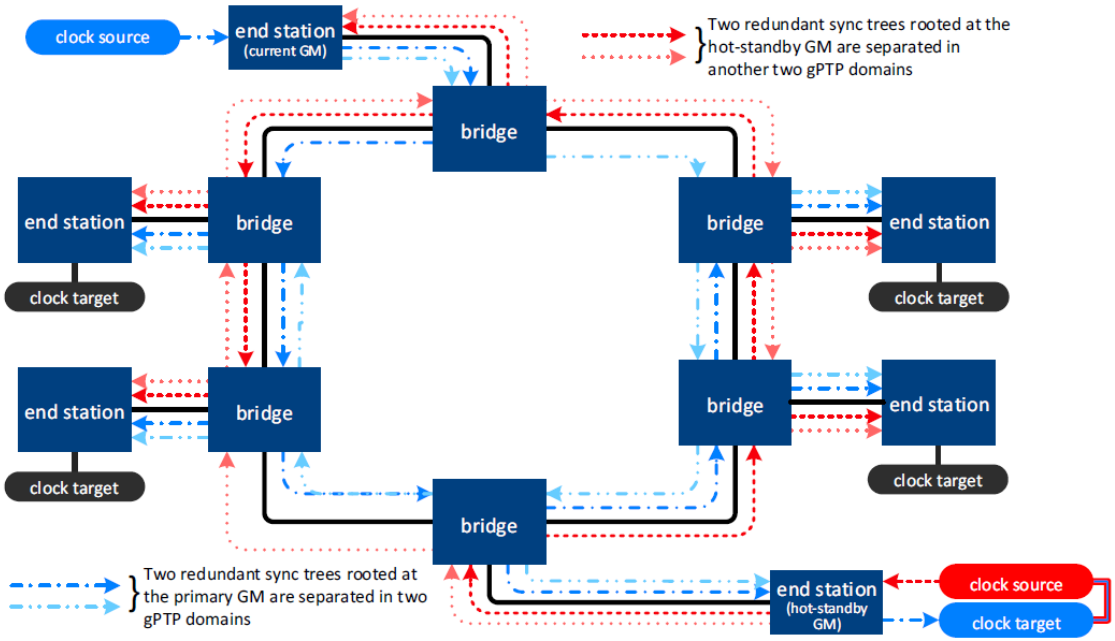

Scope to Examples: Conclusions

- Examples were designed to show what is possible
  - They serve that purpose well
- Examples do not represent practical usage
  - First two are insufficient, and third requires too much
    - P60802 performance can be met with 2 domains
  - If used for scope, the missing practicality could misdirect the project
    - Project might try to accomplish things that are not needed
- Examples do not show coupled rings

- Recommendation: Do not limit scope to examples

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 2nd approach

Scope Proposal

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

- Function to merge two or more domains into one time.
- Function to split one domain into two or more 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.
- Mitigation of byzantine faults is not supported.
- Change existing text for redundancy to align with new features.

"