A Stitch in Time

David Spencer
Product Marketing Manager, Semtech Inc.

A Stitch In Time :

- What do we mean by “a stitch in time”?
  - In the near future carriers will be able to deliver a time of day service to support applications such as LTE-A
  - But there will be other customers which cannot use this service directly
    - Imagine an industrial complex which uses a clock with its own epoch and operating rate
  - Some other potential customers already use PTP and look to the carriers to transport this over a wide area
    - There is the potential for clashes between the PTP flows of the customer and carrier
  - How can two PTP flows be stitched together to provide the needed service while avoiding clashes?

Smartgrid

- Smartgrid would be an attractive customer for the carriers, but will they be compatible?
  - Smartgrid has already selected PTP to distribute time-of-day around the power network
  - Smartgrid have developed their own PTP profile
  - Smartgrid and telecom carriers are similar
    - Same geographic coverage
    - Same customer base
    - Same financial muscle
- If they clash, who wins?

Could a Smartgrid PTP Flow be Carried Using Conventional PTP Methods?

- Could a carrier allocate a PTP domain to Smartgrid and route the flow through carrier BCs?
  - There are four problems with this approach
    - Are the BCs available for the whole route?
    - What if other customer flows need the same BCs?
    - How does the carrier maintain the flow?
    - Are the carrier's BCs compatible with the Smartgrid PTP profile?

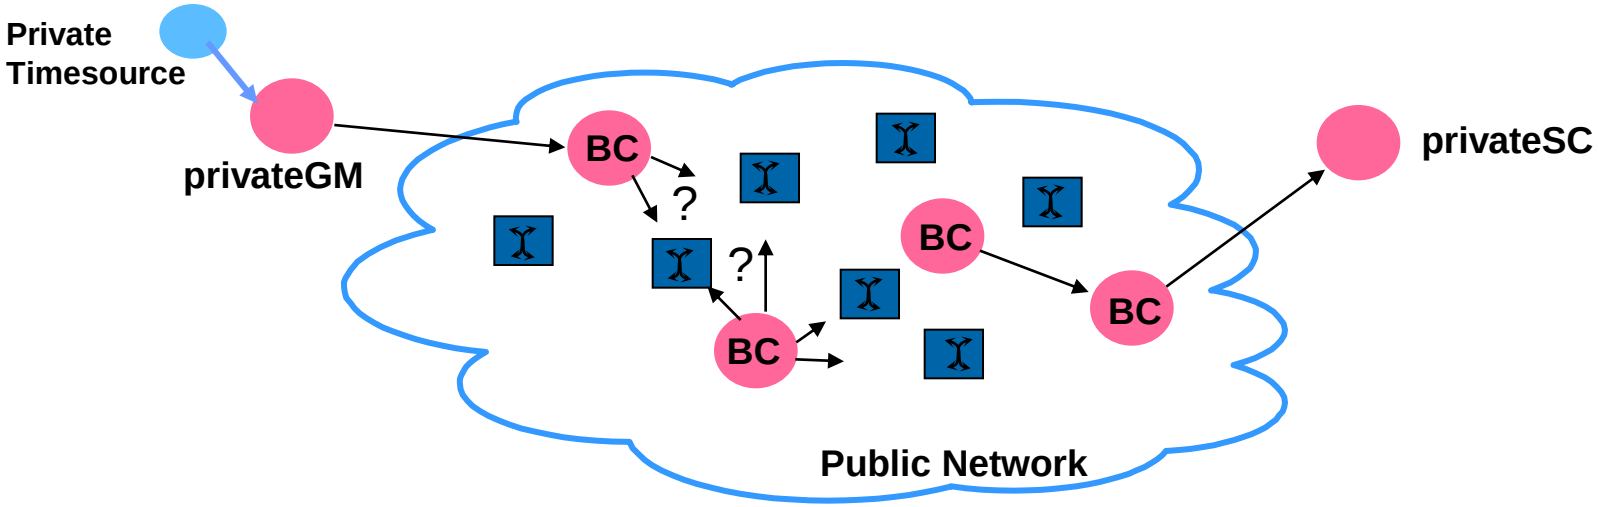

BCs Are Valuable But Limited

- Are the BCs available for the whole route?
- What if other customer flows need the same BCs?
  - IEEE 1588 says that BCs are only BCs for one PTP domain
    - PTP flows in other domains will be treated as general data flows
  - Multi-domain BCs are much more complex to implement
  - Do carriers really want to keep track of unallocated BC channels?
  - How does this technique scale?

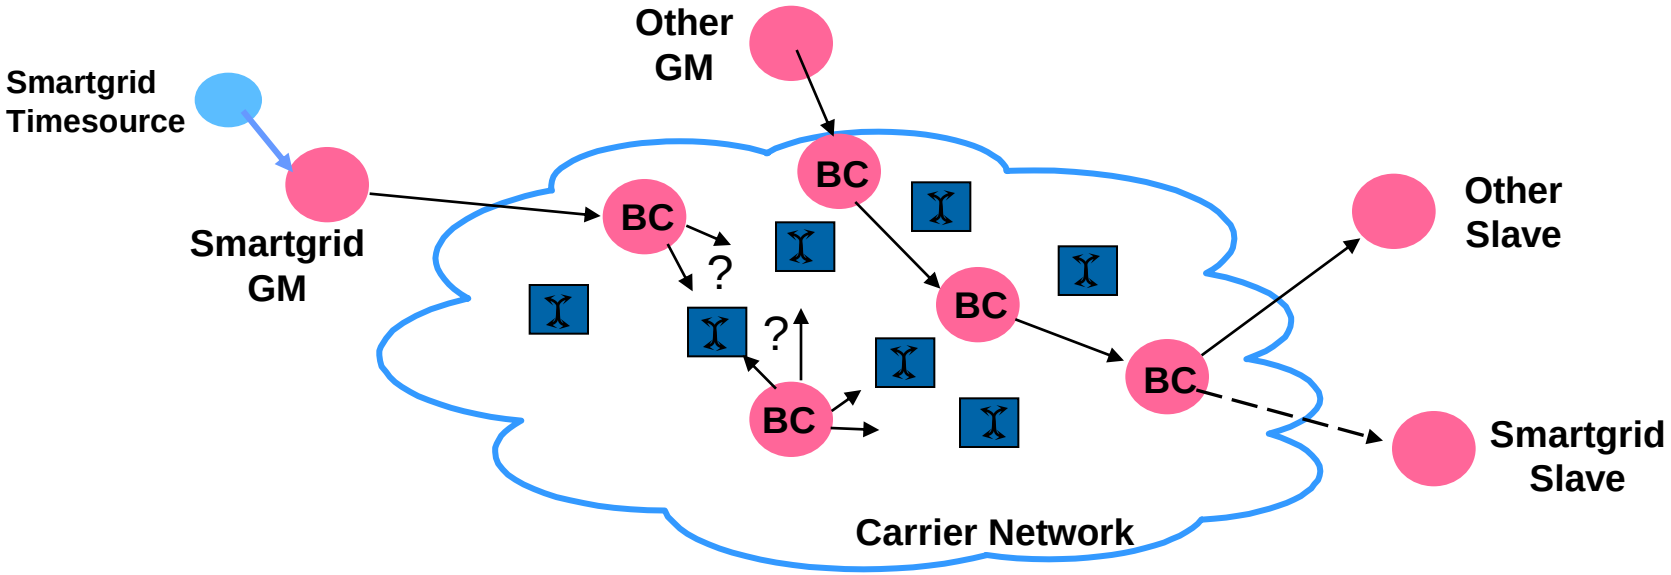

How Would a Carrier Allocate Responsibilities?

- Do BCs belong to the Sync team or the Transport team?
  - Both, of course
  - But the teams are not equivalent :
    - The transport teams are huge and used to getting their hands dirty
    - The Sync team is small and is used to precision
  - Can they work together?

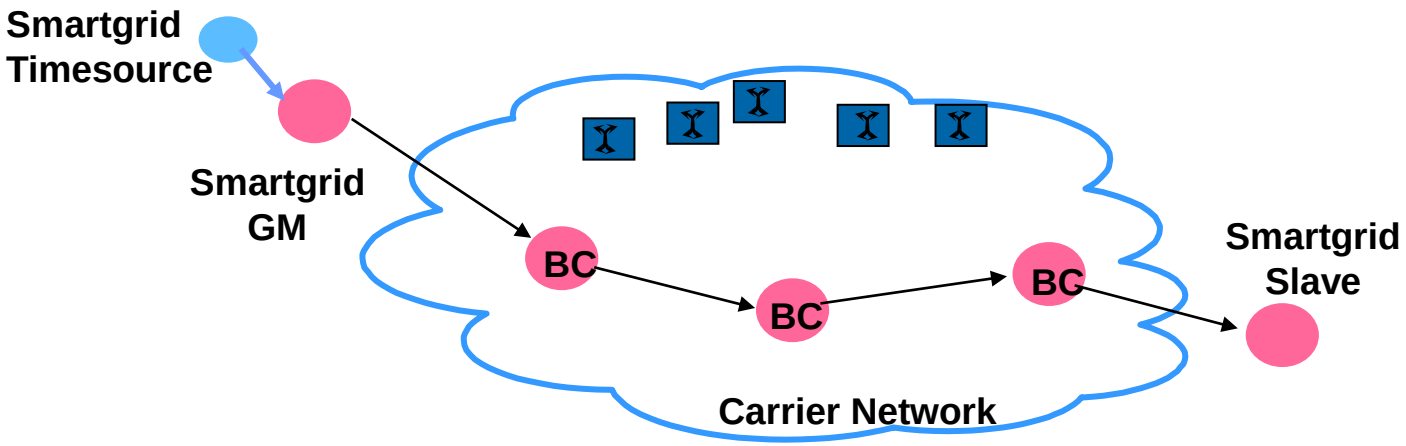

Is There a Better Way? Yes!

- **External timing could be transported through a public network as though it was just any normal form of data**
  - No on-path PTP support needed
  - But some alternative method to ensure equivalent performance
- **The Sync team can look after the carrier's timing**
- **The transport team can look after the customer flow**

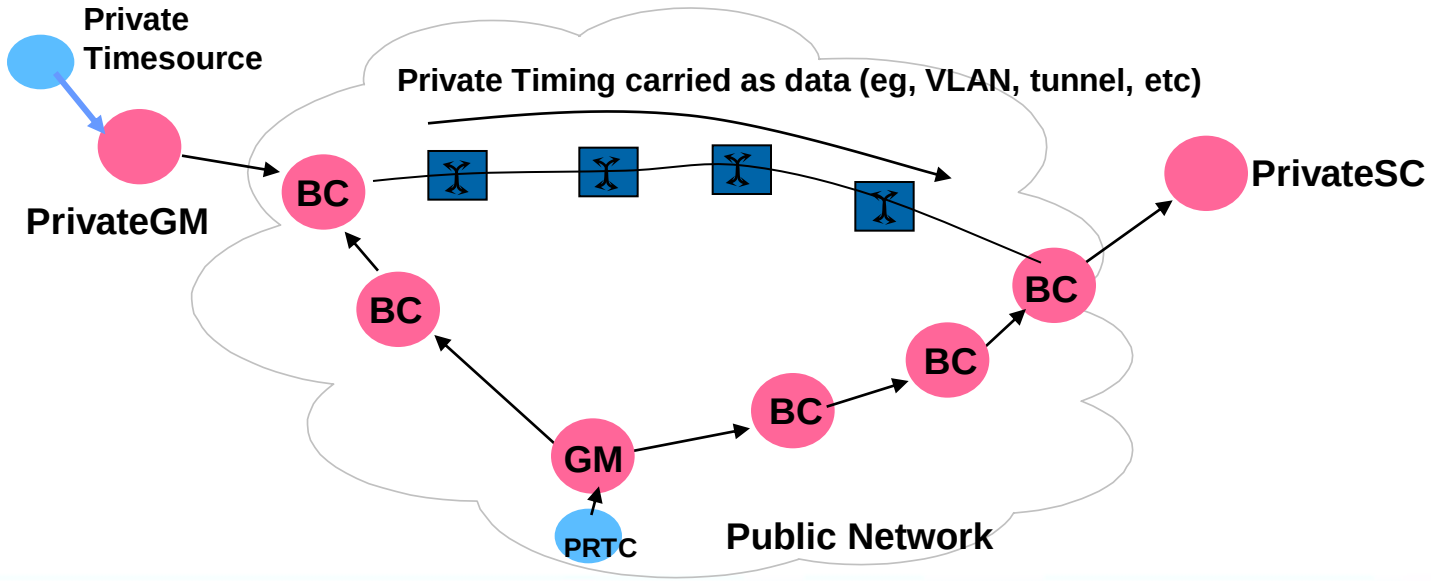

Option 1: Terminate and Recreate Private Flow

- **Three step process:**
  - Terminate incoming PTP flow at ingress and measure frequency and epoch offset
  - Transport measured offsets across network
  - Recreate PTP flow at egress using previously measured frequency and epoch offsets
- **Comparable to CES differential timing**
- **Looks simple but many details to understand**

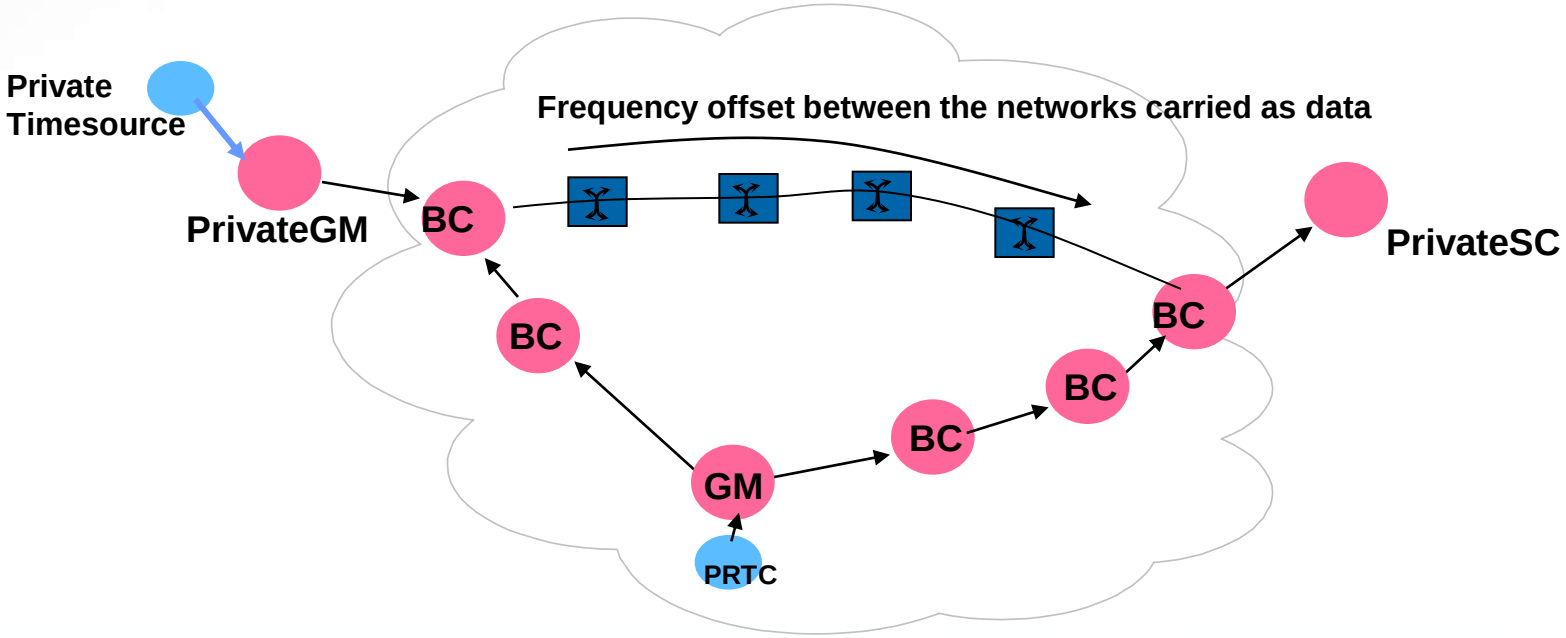

Option 1: Challenges

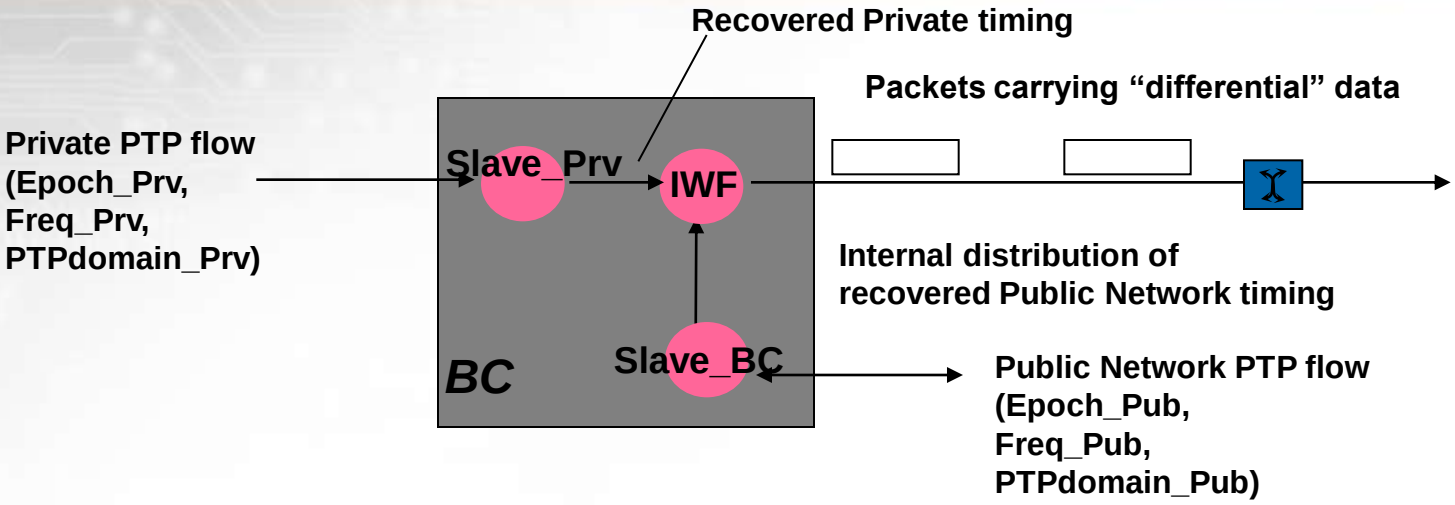

- **Ingress boundary clock would need to be slave to both carrier's timing flow and incoming private flow**
  - Would require multiple domains since BCs only have one slave port
  - Challenge 1: Care needed for domain allocation
- **Boundary clock would need to measure frequency and epoch offset between two slaves**
  - Challenge 2: How to do this robustly when network noise present

Circuit Emulation Service, According to G.8261

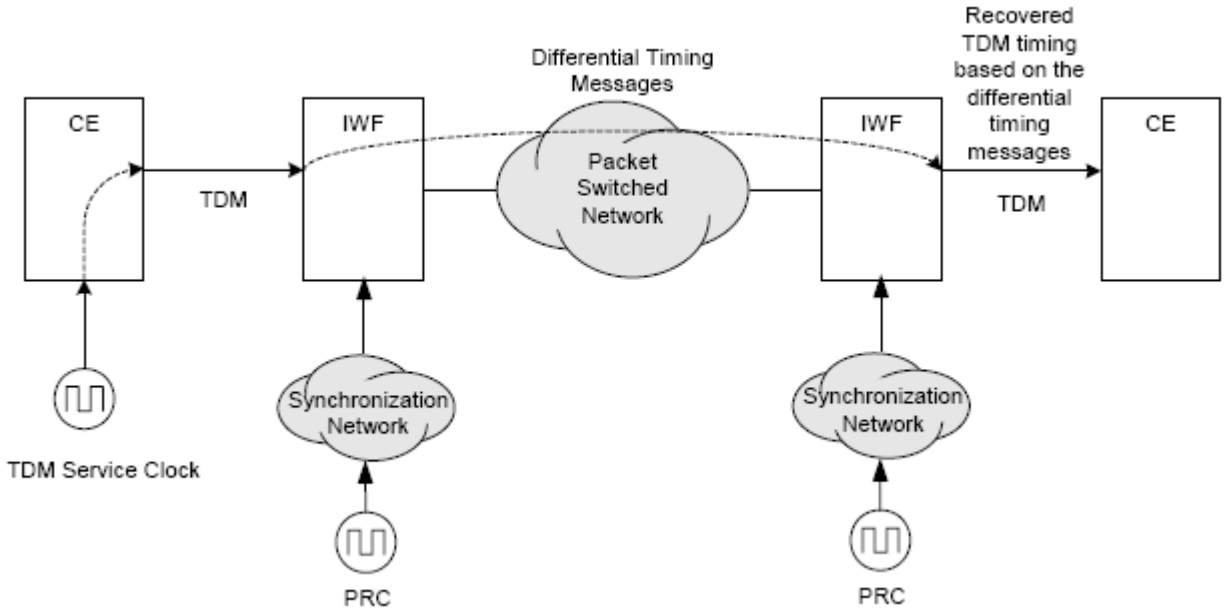

- The carrier's Sync team provide adequate timing to IWFs
- The transport team look after the customer's packet flow
- This technique scales to support thousands of customers
- We need the same for "Time-Emulation" services

Option 2: A Distributed Transparent Clock

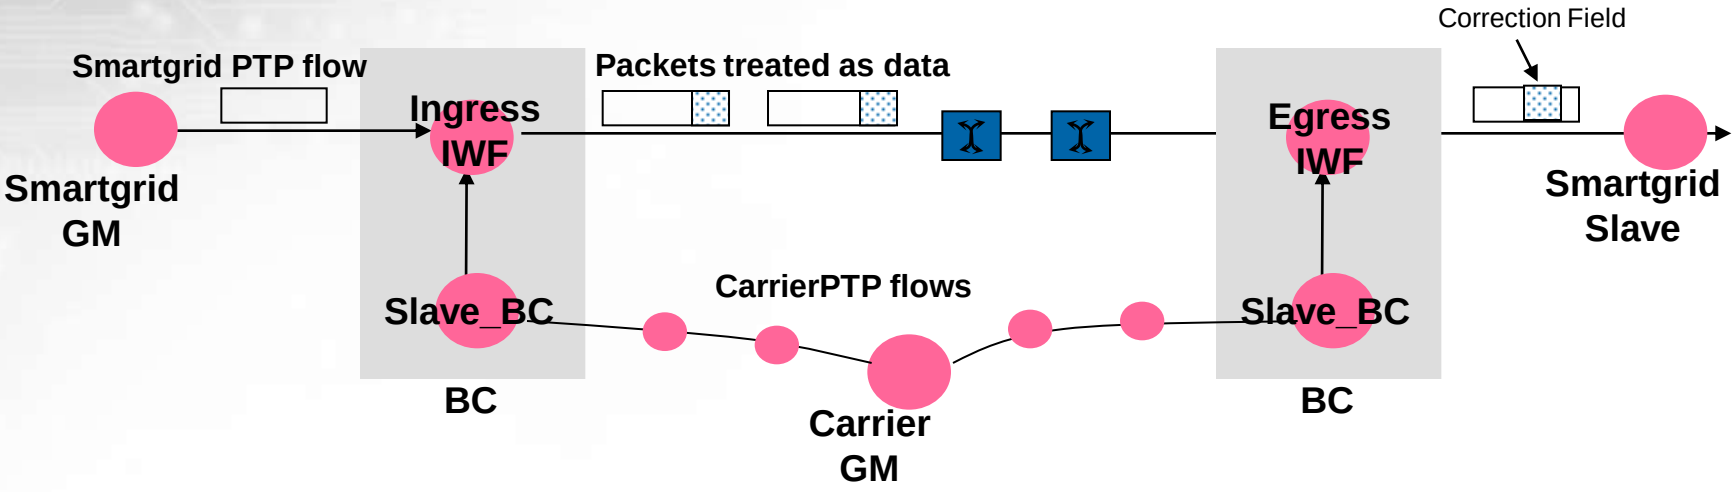

- The carrier can distribute time-of-day through interworking functions at the edge of the network
  - An ingress Interworking Function timestamps the arrival of a Smartgrid PTP packet
  - An egress Interworking Function timestamps the departure of the packet and calculates the path delay
  - This path delay can be passed to the Smartgrid Slave
- This is equivalent to circuit emulation …

A Possible Interworking Function:

- An Interworking Function would need to be bi-directional to support forward and reverse PTP message flows
  - Both sides of the IWF would timestamp the arrivals of customer PTP event messages.
  - The ingress timestamp would be sent through the carrier network by appending a timestamp TLV to the event message
  - The ingress timestamp would be subtracted from the egress timestamp to calculate the path delay of the event message
  - The path delay would be added to the correction field of event messages
- A candidate TLV has been proposed to the ITU SG15 Q13

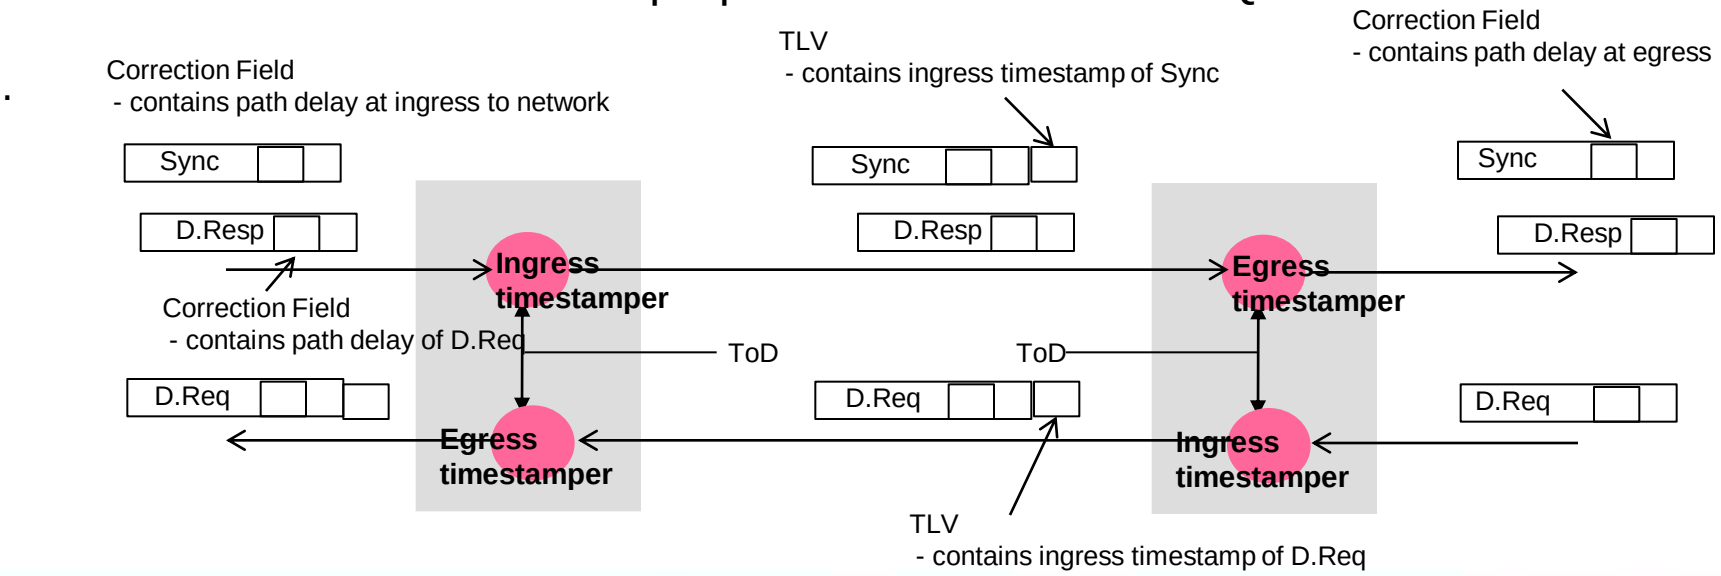

The Role of the Carrier's Sync Network:

- The distributed PTP Clock relies on the accurate delivery of time-of-day to both IWFs
  - This is the role of the carrier's Sync Network
- The distributed PTP clock works when the worst-case sum of errors at each end is within the needs of the customer network
- Carriers are adopting BCs and Synchronous Ethernet to provide good accuracy at the edge of the network.
- What level of performance can be expected?

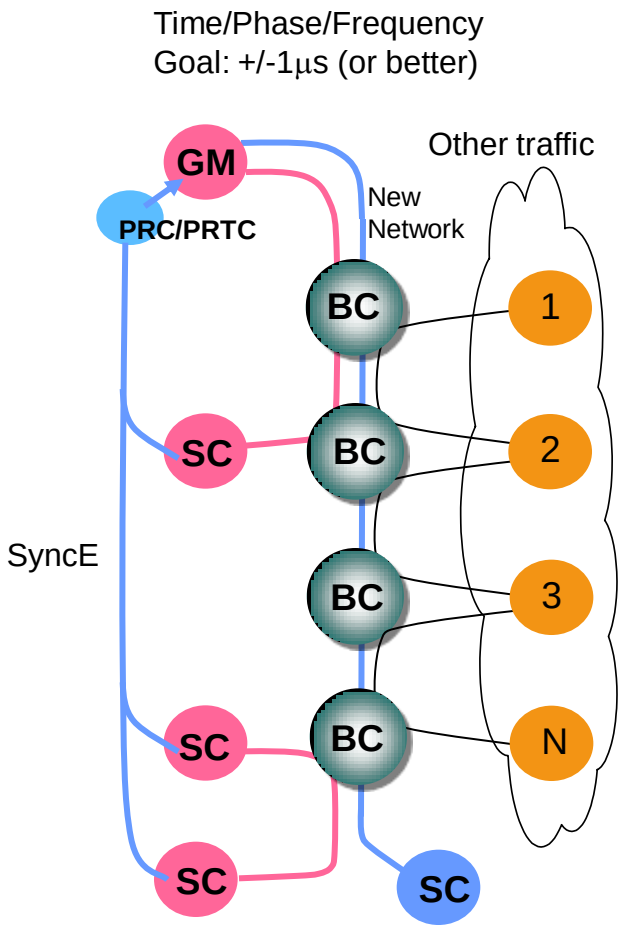

G.827x - Telecom Profile 2

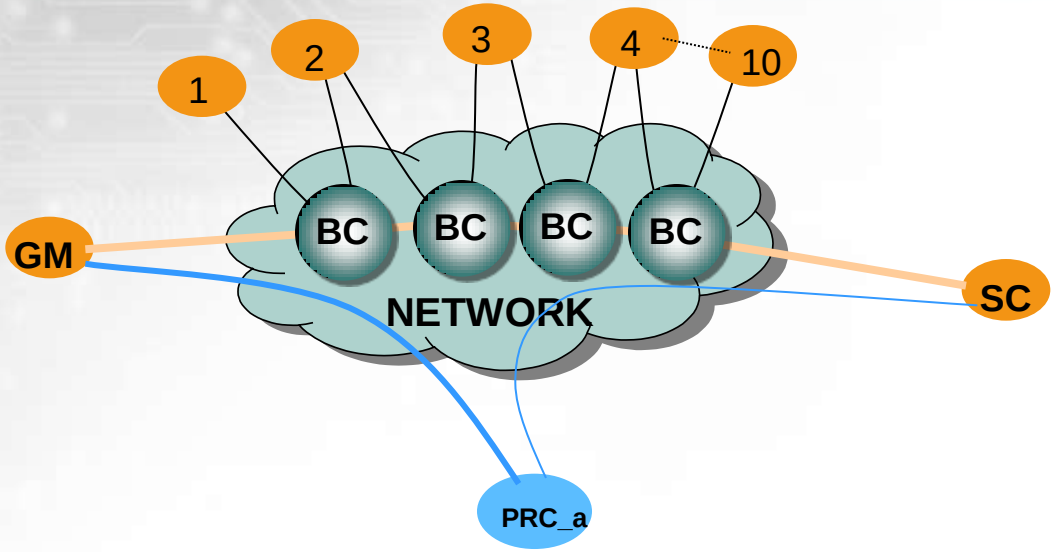

Coherent: Common source of frequency used at each end

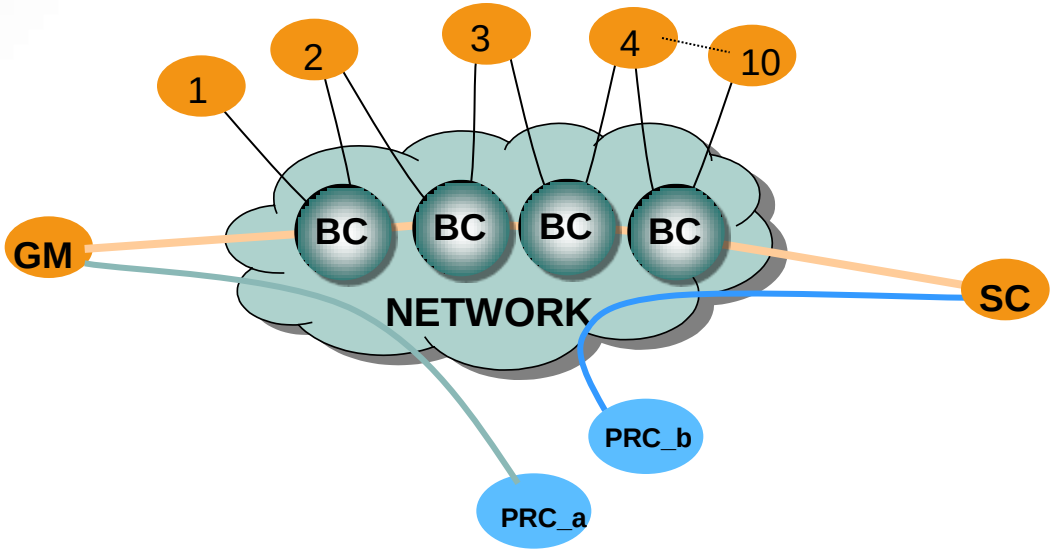

Congruent: Different sources of frequency used at each end

Performance Under Failure Mode

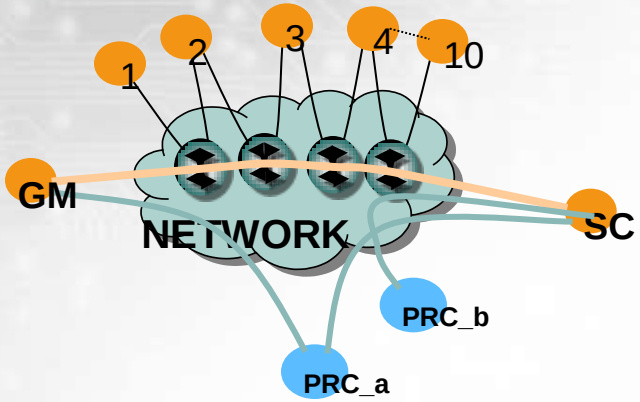

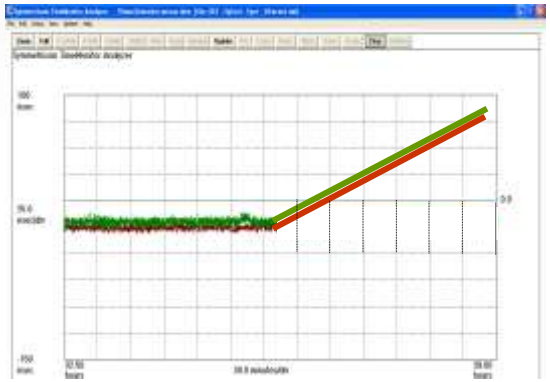

In this example, the Slave Clock used two key features:
1 - the PHY clock stabilising the Slave was initially traceable to the same reference clock that is driving the GrandMaster (Co-herent timing) .
2 - the PDV-processing algorithm was strongly biassed towards using early detection of minimally-delayed packets, one in each direction.

With this combination the Slave Clock shows good phase/time accuracy and extraordinary phase stability. This is because, at the beginning of the test, the algorithm detected a suitable pair of minimally-delayed packets which it used to get good alignment of the Slave Clock. Good alignment was then maintained by driving the Slave at the same rate as the GrandMaster. Traffic-induced PDV is reduced by more-or-less ignoring subsequent packet delays. The output noise is essentially that of the physical-layer clock. This method has been proposed by several contributors.

While the technique looks promising, the Slave is heavily dependent on the performance of the network clock. Excessive noise, phase transients and frequency deviations would quickly spoil the Slave Clock output. The TIE plot to the left shows what could happen when the sync network has to switch to an alternate source.

The following slide shows a modification of the technique aimed at improving the Slave Clock's robustness.

Can Legacy Equipment be Upgraded To Deliver Similar Performance? Yes, Using pBCs…

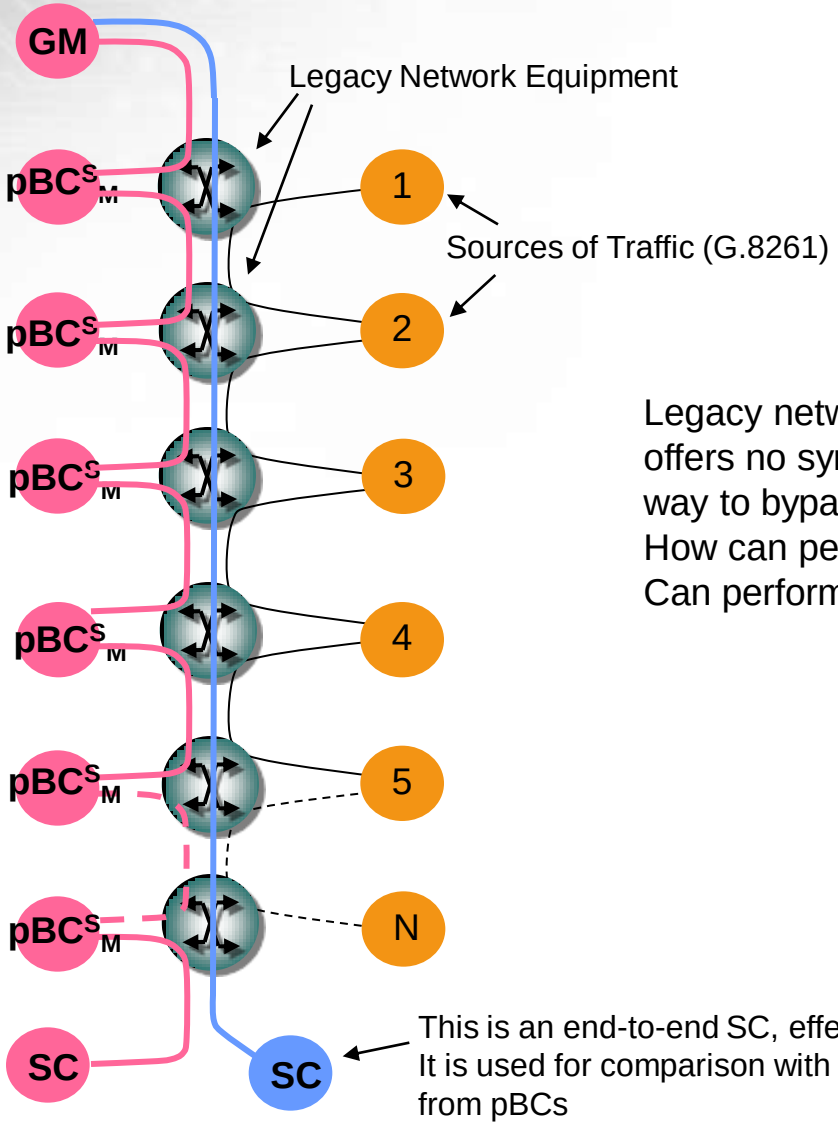

Plug-in Boundary Clock (p**BC)**

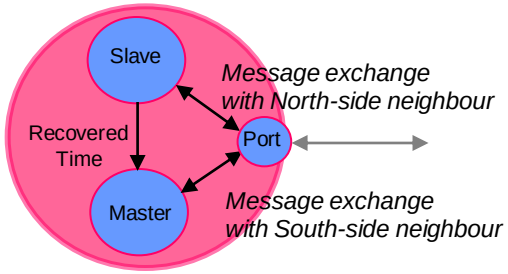

Legacy networks have been built from equipment which offers no syntonisation from a PHY-layer clock and no way to bypass PDV-generating functions.
How can performance be improved?
Can performance approach Telecom Profile 2?

This is an end-to-end SC, effectively running under Telecom Profile 1.
It is used for comparison with the performance of a Slave running with support from pBCs

Phase Deviation With 10-node pBC Chain

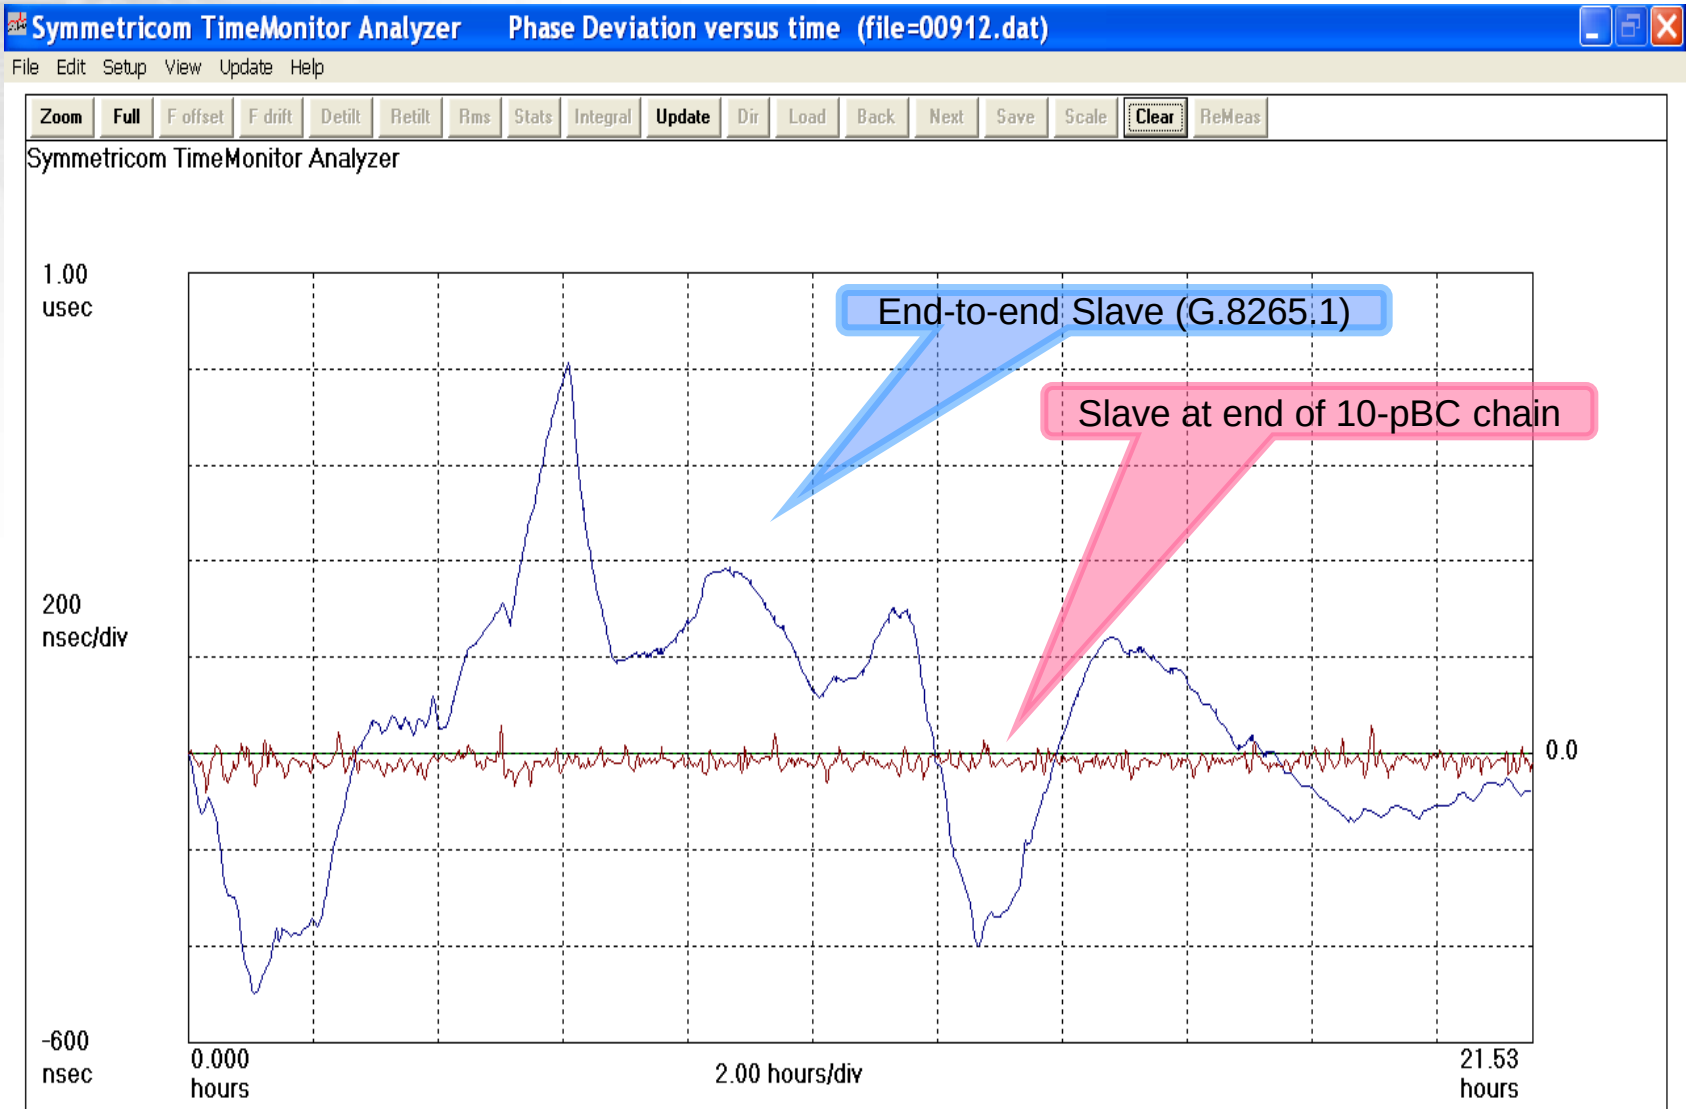

Phase Deviation With 10-node pBC Chain

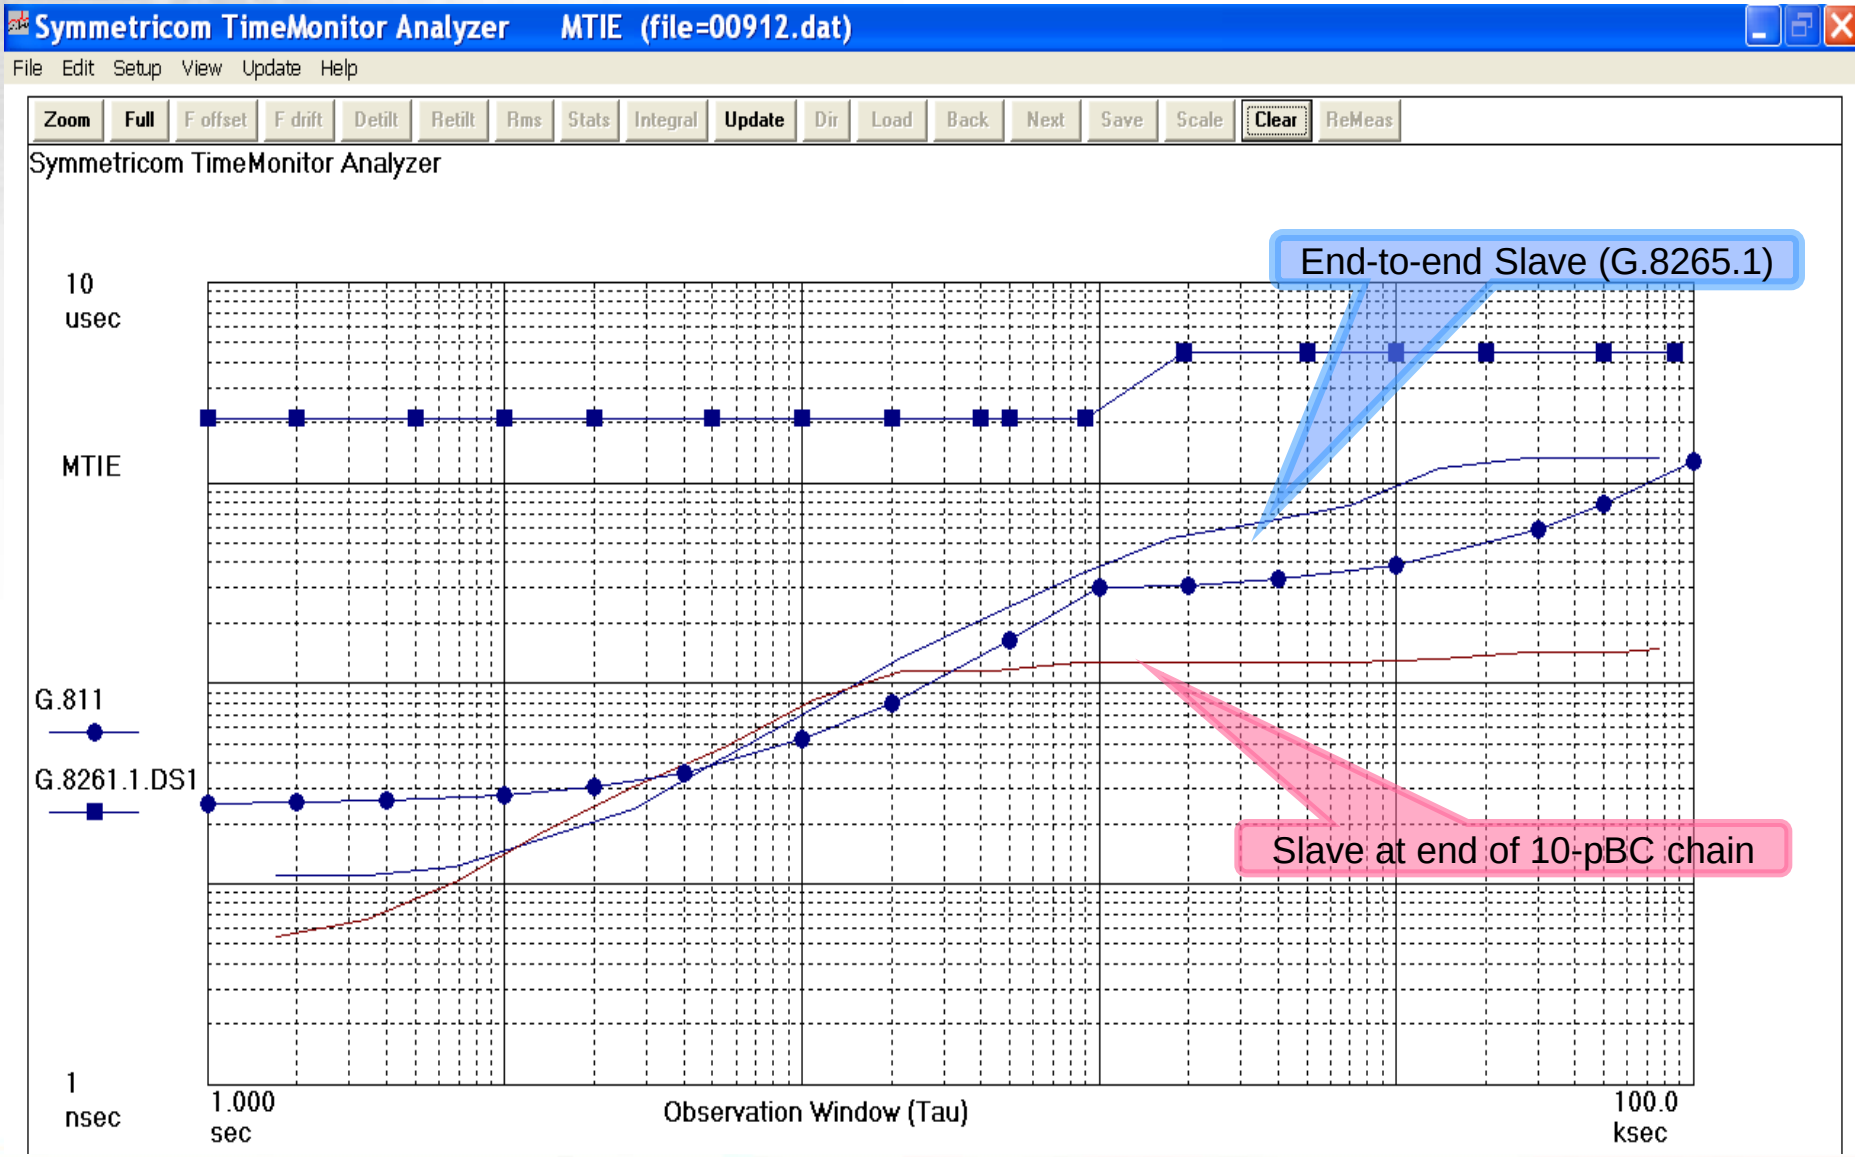

What Would a Plug-in Boundary Clock Look Like?

A plug-in BC exchanges timing packets with neighbours in both directions.
It supports two flows of timing packets but only requires one Ethernet port on a legacy switch.
It could be built as an add-on box or even, in principle, as a small-form plug-in module.

Legacy Ethernet switch
(eg, not PTP-aware, no SyncE support)

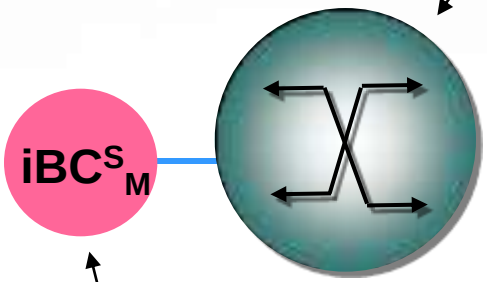

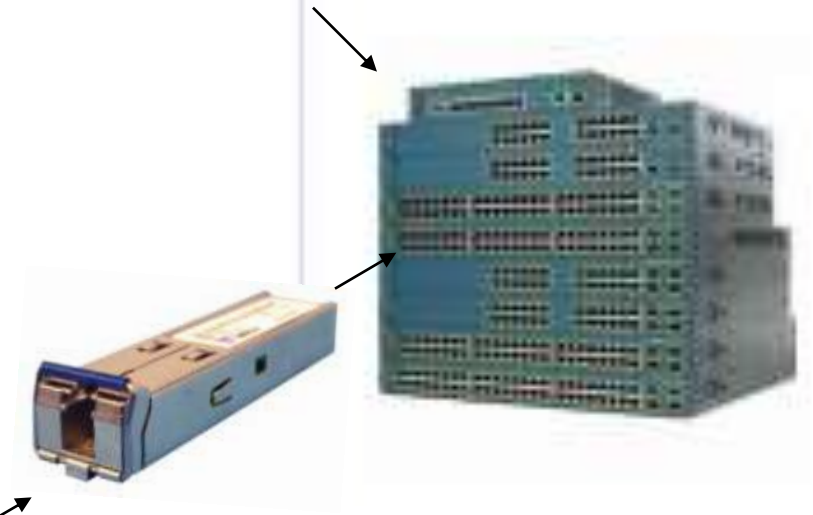

?

Interim Boundary Clock
(i.e., combined Slave and Master)

A Plug-in BC is not an alternative to G.827x BC/SyncE, but just a way to upgrade a legacy ethernet switch where useful

Summary

- Future applications will likely require transport of multiple private timing flows over carrier networks
  - Independent of carriers own sync mechanism
  - Could be in parallel with “sync as a value-added service”
- This can be achieved by using distributed transparent clocks
  - Uses inter-working function at ingress and egress ports
  - Separates sync and data transport responsibilities
- Good synchronization within the carrier’s network will still be essential
  - Likely to require universal use of boundary clocks
  - Plug-in boundary clocks provide an alternate solution
    - Upgrade of legacy networks
    - Provision of low-cost embedded solution
    - Semtech has demonstrated sub-100ns performance over 30 hops without SyncE support