Practical Use Cases for Ethernet Redundancy

Don Pannell, Fellow
Automotive Ethernet Networking, NXP Semiconductor

Nicolas Navet, Professor
University of Luxembourg / Cognifyer.ai

IEEE Ethernet & IP Tech Days – September 2020

SECURE CONNECTIONS FOR A SMARTER WORLD

OVERVIEW

- Examined Redundant Network Configurations
- Soft Error Rate Modeling
- Hard Error & Cost Factor Considerations
- Summary

THE NEED

Zonal

- Zonal networks, as shown in the figure, easily support redundancy, especially on the highlighted Ethernet backbone
- IEEE 802.1CB is the TSN standard for Seamless Redundancy, supporting zero recovery time from lost frames
- This presentation models CB in various topologies with the CB function at various locations in the network so these differences can be quantified
- This is NOT a Safety presentation, but the data presented will help Network & Safety teams develop cost effective redundancy solutions

Examined Redundant Network Configurations

Single changes to a baseline are compared so their effect can be extrapolated

SECURE CONNECTIONS
FOR A SMARTER WORLD

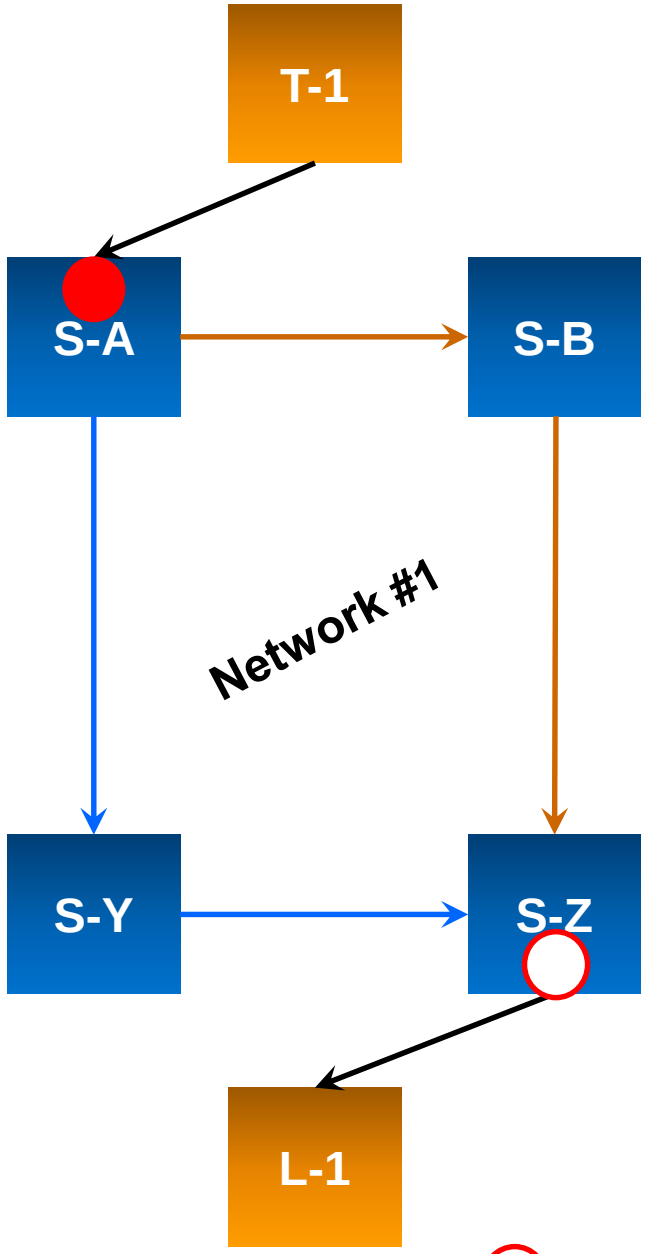

4 CB Switches, Non-CB Talker & Listener

- Pros:
  - Existing end-nodes can be used
  - Hard or Soft errors on the backbone wires are protected
    - A Hard error is a long duration error like a broken wire
    - A Soft error is an intermittent error like a CRC errored frame
    - The Brown & Blue wires are redundant paths for the packets
  - Failure of S-B or S-Y is protected
- Cons:
  - Backbone bandwidth is double+ for the redundant flows
    - The '+' is due to the added 6-byte R-Tag & possible 4-byte S-Tag
  - Links from the Talker & to the Listener are not protected
    - The Black wires
  - Failure of S-A, T-1, S-Z or L-1 is not protected

= CB Seq # & Split = CB Merge

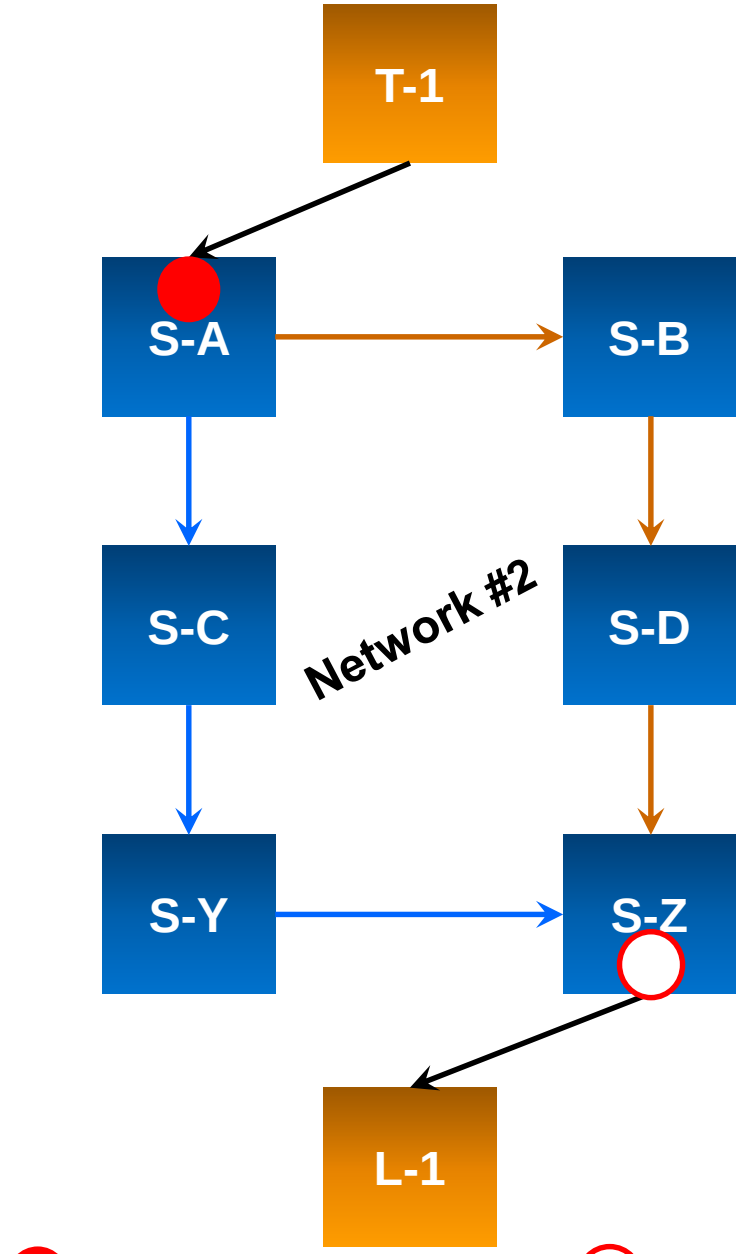

6 CB Switches, Non-CB Talker & Listener

- Very similar Pros & Cons as the previous slide
- Only change is the addition of S-C & S-D in the backbone
- Modeled to see the impact of Soft errors on the increased number of links in the protected backbone

= CB Seq # & Split = CB Merge

Non-CB Switches, CB Dual-Homed Talker & Listener

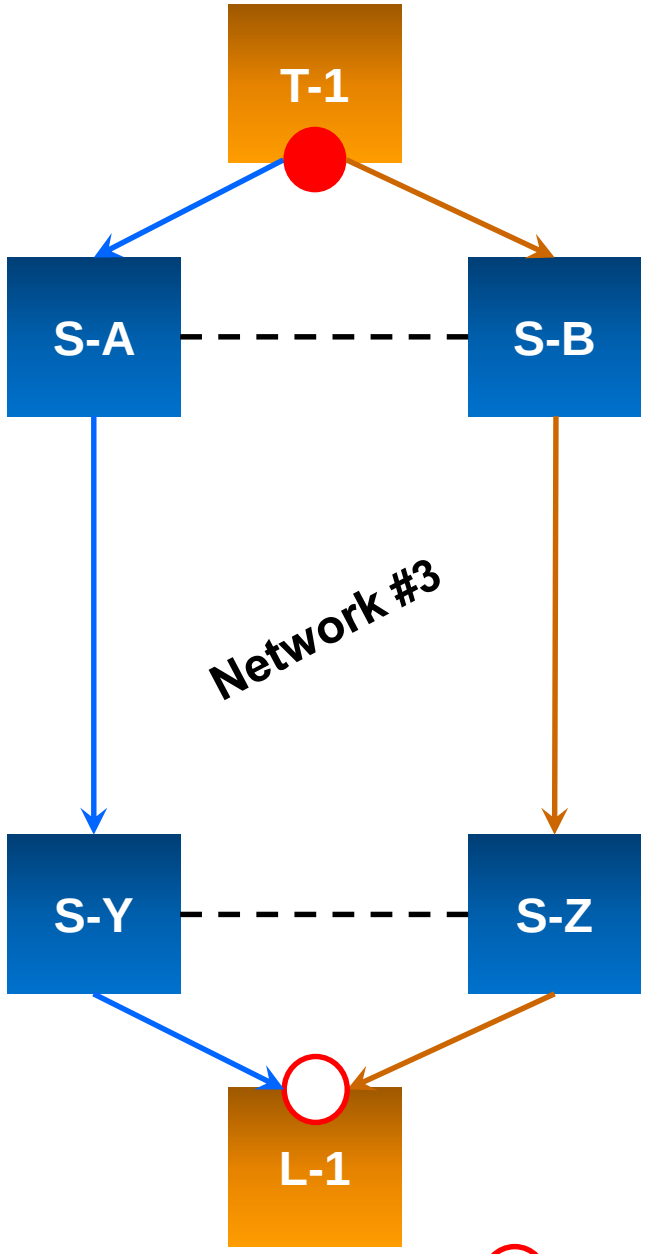

- Pros:
  - Existing switches can be used
  - Hard & Soft errors on the entire path are protected
    - The Brown & Blue wires (dotted wires are not used for these flows)
  - Failure of any single switch is protected
  - Failure of anything in the Blue path (wire or switch) is protected
  - Failure of anything in the Brown path is protected
  - Backbone bandwidth is half of Network #1 for the redundant flows as dotted wires are not used & are available
- Cons:
  - Requires dual-homed end nodes (with dual Ethernet ports)
  - End nodes replicate frames & eliminate the duplicates
  - Failure of T-1 or L-1 is not protected

● = CB Seq # & Split ○ = CB Merge

Non-CB Switches, CB Single-Homed Talker & Listener

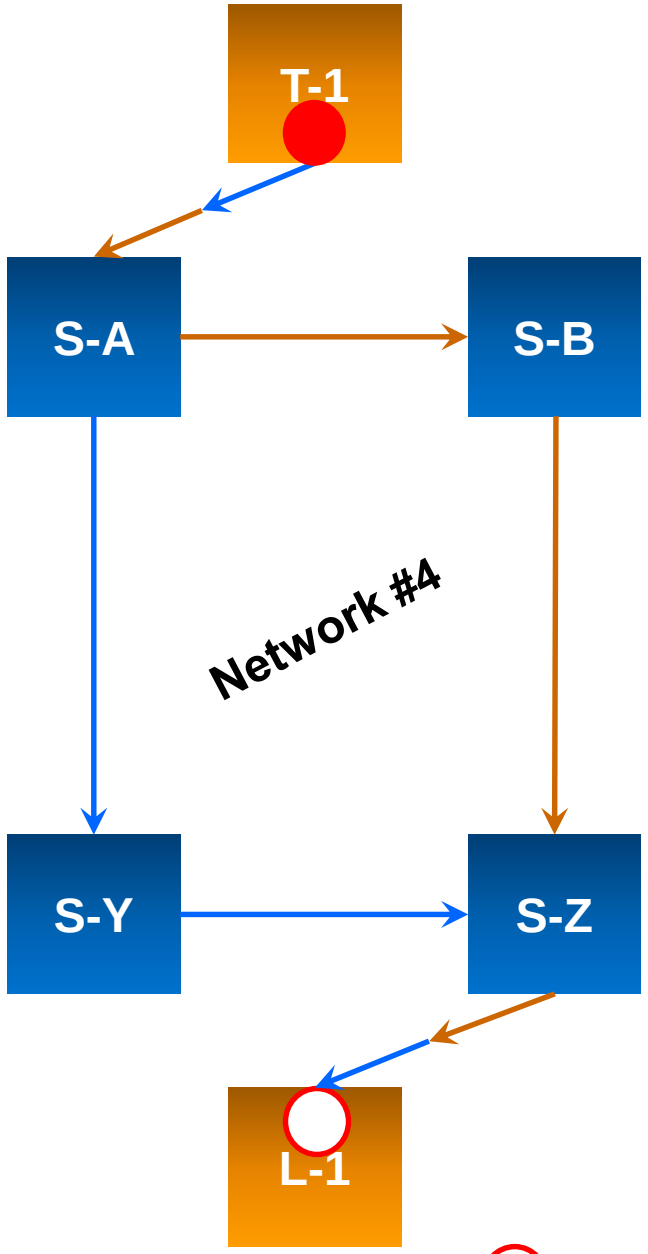

- Pros:
  - Existing switches & end nodes w/new software can be used
  - Hard or Soft errors on the backbone wires are protected
    - The Brown & Blue wires between the switches
  - Soft errors on the links from the Talker & to the Listener are protected
    - Temporally due to the doubled transmission
  - Failure of S-B or S-Y is protected
- Cons:
  - End-to-end bandwidth is double+ for the redundant flows
    - The '+' is due to the added 6-byte R-Tag & possible 4-byte S-Tag
  - End nodes replicate frames & eliminate the duplicates
  - Failure of S-A, T-1, S-Z or L-1 is not protected

● = CB Seq # & Split ○ = CB Merge

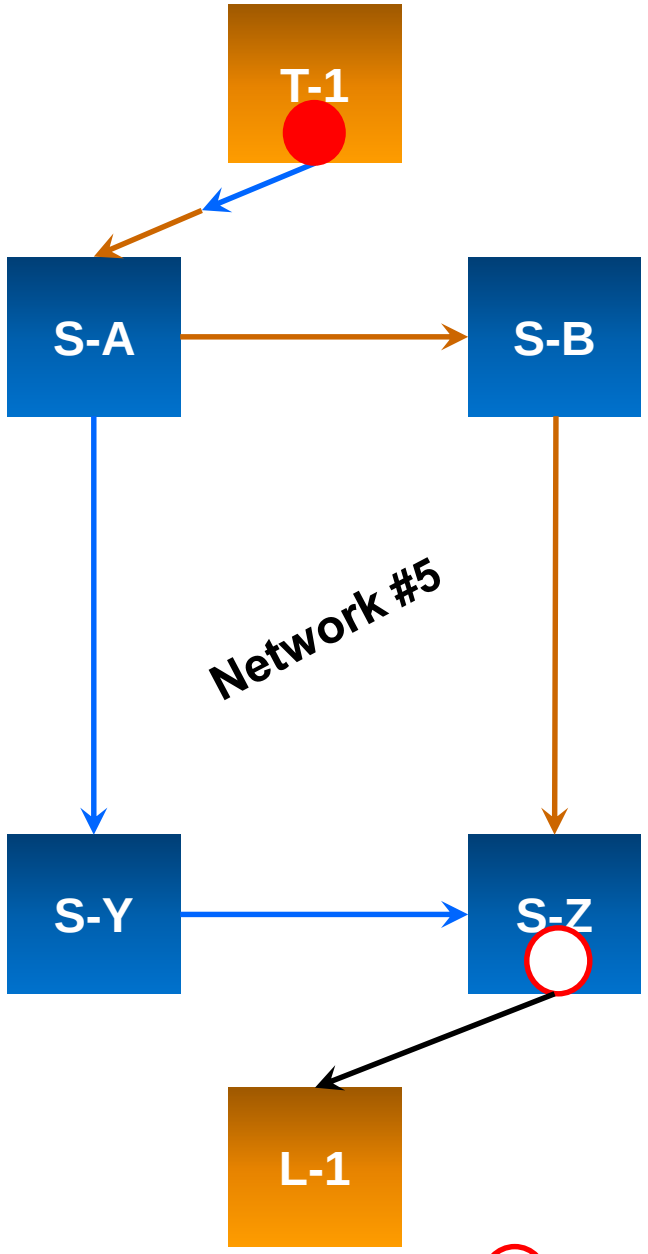

Mixed Switches, Mixed Talker & Listener

- Very similar Pros & Cons as previous slide
- Only change is S-Z does the duplicate frame removal
  - This means the S-Z to L-1 link is no longer protected
  - But existing ECU's without any software changes can be used
    - With increased redundancy support achieved via firmware updates
    - i.e., Network #1 → Network #5 → Network #4
    - CB enabled switches are needed in this case
- Modeled to see the impact of Soft errors on the unprotected link to the Listener
- This mixture is supported as long as T-1 creates the Brown & Blue frames in accordance to 802.1CB

= CB Seq # & Split = CB Merge

Soft Error Rate Modeling

End-to-end protection is key for high-integrity communication

SECURE CONNECTIONS
FOR A SMARTER WORLD

Error model used for soft errors

- Soft errors are limited to frames dropped due to CRC errors
- Bit Error Rate (BER) is assumed identical on all links & constant over time
- CRC errors are independent of each other, i.e., no “bursts” of errors
- 100BASE-T1 specifies BER ≤ 10^{-10}, PHYs are much better in practice, thus a BER of 10^{-12} is used in the experiments

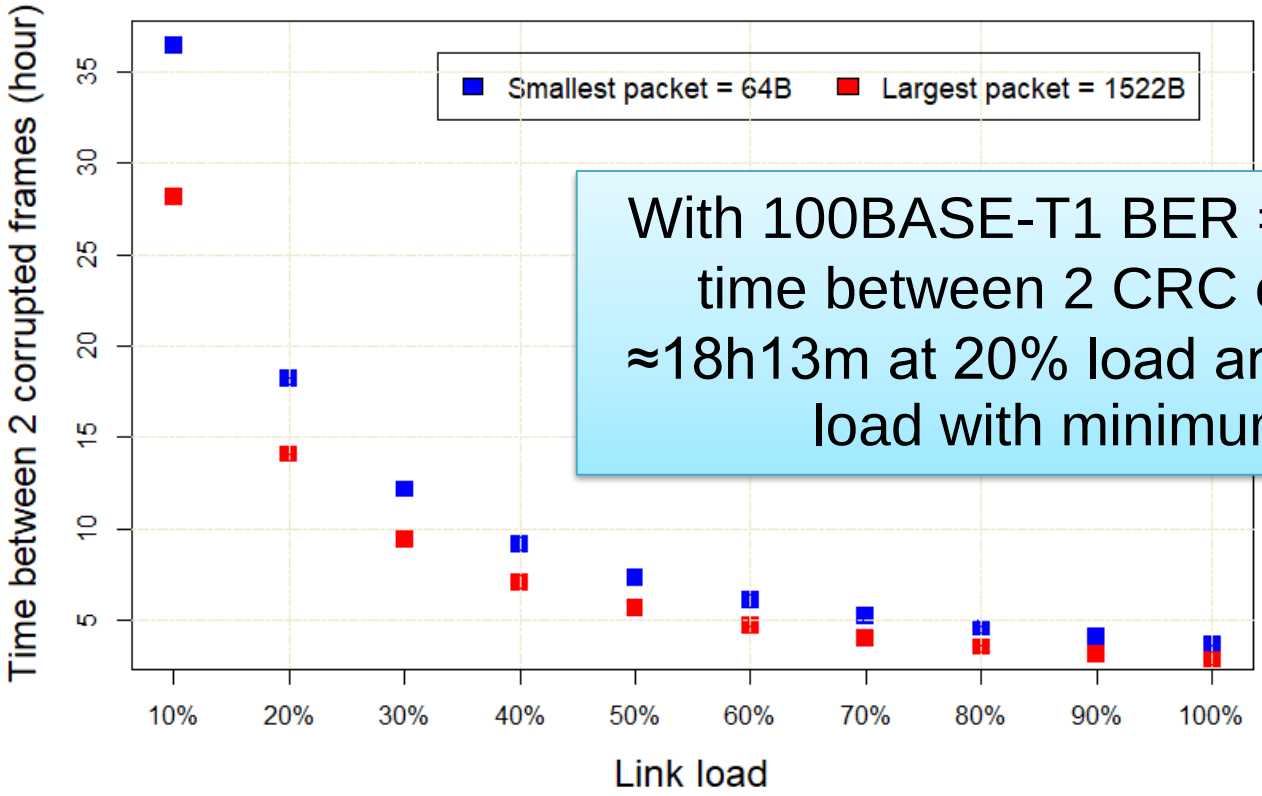

With 100BASE-T1 BER = 10^{-12}, the average time between 2 CRC errors on a link is ≈18h13m at 20% load and ≈3h38m at 100% load with minimum frame size

Replication solutions : requirements & Single Points of Failure

2 CB-aware switches

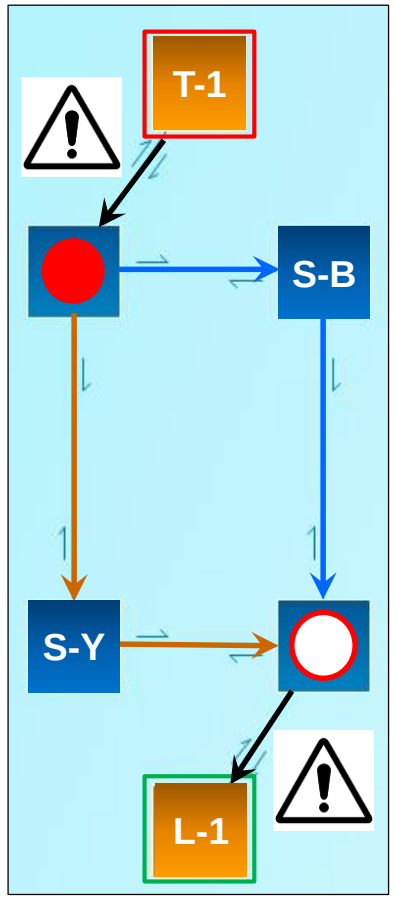

1

2 CB-aware switches w/more hops

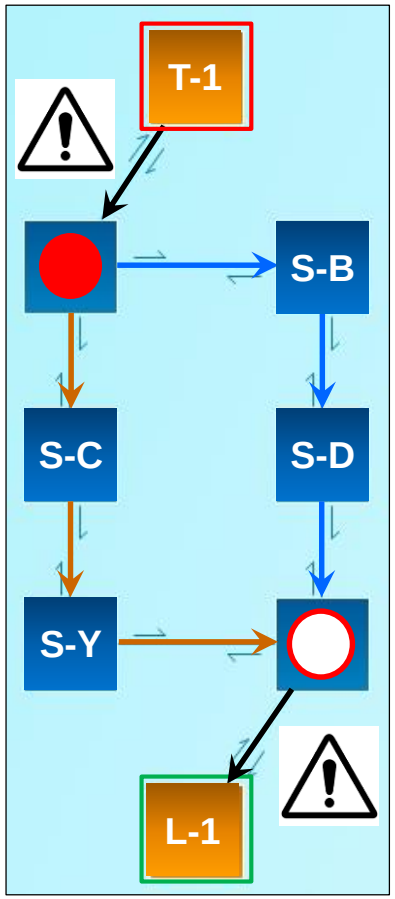

2

CB-aware talker & listener + dual-homing

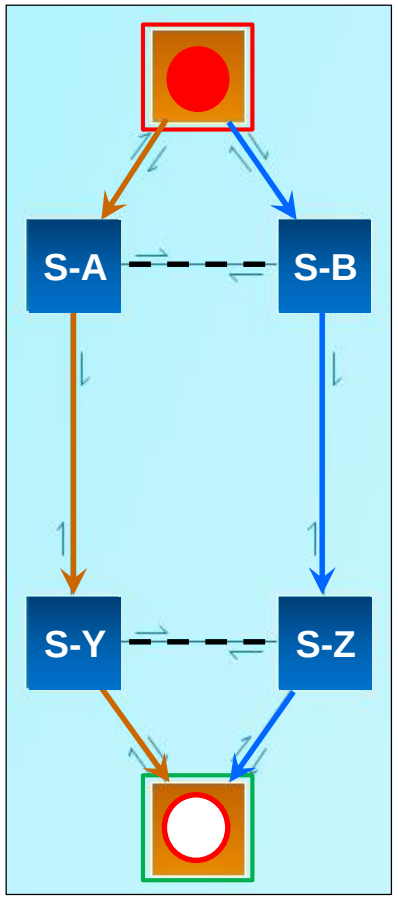

3

CB-aware talker & listener + temporal redundancy

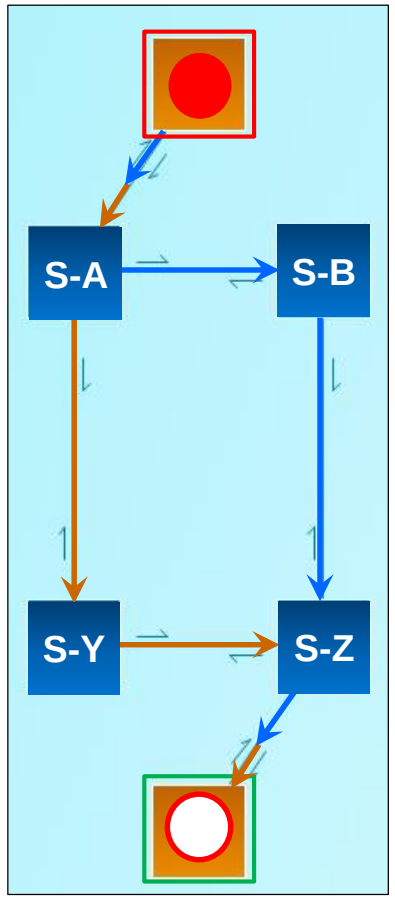

4

CB-aware talker + 1 temporal redundancy + 1 CB-aware switch

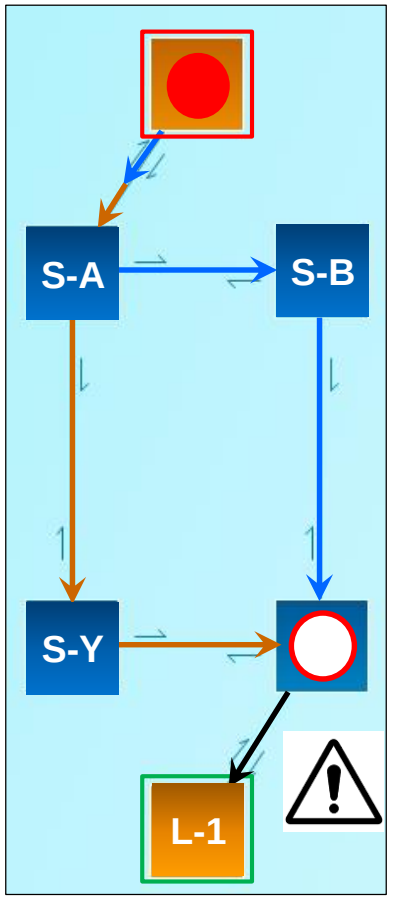

5

[RTaW-Pegase screenshots]

= CB Seq # & Split

= CB Merge

= unprotected transmissions on that link → Single Point of Failure for Soft Errors

Replication in action - packet loss rate

- Packets are lost when none of the copies are received by the listener(s)
- The data assumes a homogeneous Bit Error Rate = 10^{-12} regardless of link speed
- 2.44E-8 means 2.44×10^{-8}

Network	Loss Rate for 1522B Packet	Loss Rate for 64B Packet	Improvement factor vs. No Redundancy	Takeaways
1	2.44E-8	1.02E-9	2	The two unprotected transmissions are by far the dominant risk factor
2	2.44E-8	1.02E-9	2.5	# of hops is a low order factor as long as transmissions are protected
3	1.33E-15	2.36E-18	8.7E8	Very robust to soft errors!
4	2.37E-15	4.19E-18	4.9E8	Same order of magnitude as #3
5	1.22E-8	5.12E-10	4	Twice as robust compared to #1 due to the single unprotected transmission link

3 (dual-homing) and 4 (end-to-end temporal redundancy) stands out
3 additionally protects against any one hard error unlike 4

Replication in action – average time between 2 packet losses

- Assuming a transmission period of 1ms with min. size frames (e.g., actuator messages)
- The data assumes a homogeneous Bit Error Rate regardless of link speed
- Stated times are for one flow only with the low link utilization as stated above

Network	BER = 10^{-11}	BER = 10^{-12}	BER = 10^{-13}	BER = 10^{-14}
Baseline: no redundancy	13hours 33min	5days 15hours	56days 12hours	≈ 1year 6months
1	1day 3hours	11days 7hours	113days	≈ 3years 1month
2	1day 3hours	11days 7hours	113days	≈ 3years 1month
3	1.3E5 years	1.3E7 years	1.3E9 years	1.3E11 years
4	7.5E4 years	7.5E6 years	7.5E8 years	7.5E10 years
5	2days 6hours	22days 14hours	226days	≈ 6years 2months

Hard Error & Cost Factor Considerations

Nothing is free, and not all solutions cost the same!

Cost Factor – Components

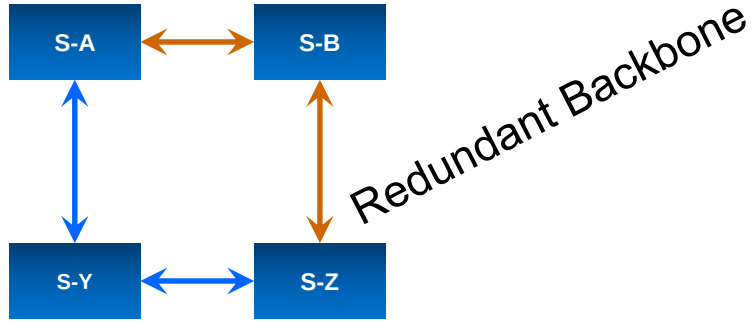

- Creating a Redundant backbone:
  - 1st: Create a Ring network from a daisy-chain one
    - Cost = 1 extra link only in the network
  - 2nd: Add seamless redundancy: 802.1CB support in bridges in the critical data's path
    - Doubles the backbone's bandwidth for the critical flows
    - Cost = varies depending on the CB requirements needed, like the bandwidth & the number of critical flows

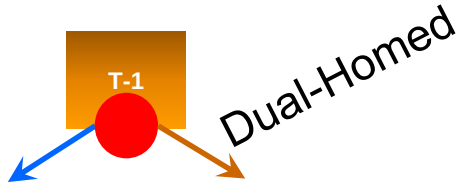

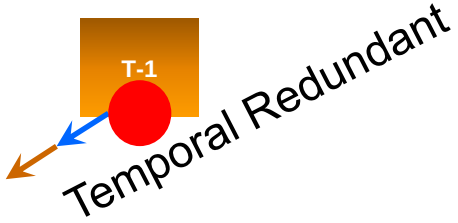

- Redundancy in the first & last links:
  - Option A: Use dual-homed end nodes
    - May keep the backbone's bandwidth close to the same loading as before
    - Cost = 1 extra link per critical end node & more CPU cycles to run 802.1CB for the critical flows only
  - Option B: Use temporal redundancy single-home end nodes saving the cost of the extra links
    - But the backbone's bandwidth is still doubled and it also adds…
    - Cost = More CPU cycles to run 802.1CB & duplicate frame transmission for the critical flows only

Cost Factor – For Sensors & Their Flows

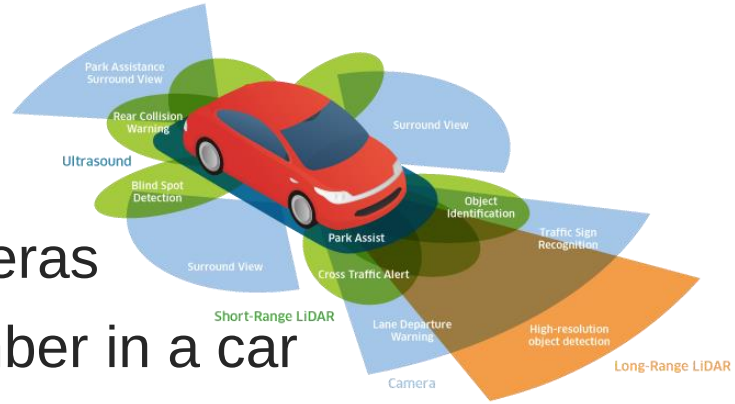

- Sensors can be very high bandwidth devices like 8+ gig/sec cameras
- Many sensors are needed too, 8 cameras is becoming a low number in a car
- Can this data even fit on a backbone today (even with reduced requirements)?
  - With 6 camera @ 4 gig/sec requires 24 gig/sec! With redundancy that grows to 48 gig/sec.
  - What year will these Ethernet PHYs & Switches be cost effectively available for Automotive?
- This bandwidth does not take into account the added data for Lidar, Radar, etc.
- And why double the bandwidth requirements of a sensor when the sensor itself is a single point of failure?
  - A camera “failure” is more likely due to dirt on the lens vs. a silicon or a wire failure!
  - Why not add more cameras such that they overlap instead? Then CB is not needed for these flows
- Sensor Fusion merges data from many cameras & other sensor types to form a “picture”
  - This process repeats continuously such that a Soft error is likely not critical and even many Hard errors (like dirt) may allow for continued operation at reduced speeds (i.e., limp home)

Cost Factor – For Actuators & Their Flows

- After Sensor Fusion, a decision is made on what to do
  - Steer away from a problem or slow down, etc.
- Turning, braking, accelerating, etc. are Actuators

- Actuators:
  - Are very low bandwidth devices
    - Historically Actuators have been connected using CAN and LIN
    - Thus doubling the bandwidth of these flows on the Ethernet backbone is totally feasible
    - Even doubling the bandwidth, temporally, on a single link to the backbone is feasible
  - Re-transmission (temporal redundancy) is quite often used today to overcome Soft errors
  - Seamless redundancy greatly increases the likely hood that the 1st transmission is received
    - Thus keeping the latency time for these Decision to Action messages low

Summary & Conclusions

Summary & Conclusion

- Hard, persistent errors are easier to see & plan around as noted in the examples
  - The likely hood & impact of these errors (wire vs. software vs. silicon) is application dependent
- Soft, intermittent errors are harder to evaluate, but the impact of BER has been shown
  - 1000BASE-T1 supports FEC making BER analysis much harder (possible future work?)
  - The numbers presented for 100BASE-T1 are felt to be a good rule-of-thumb for 1000BASE-T1
- Redundancy is not free & the biggest cost today could be the extra backbone bandwidth
  - Therefore, apply redundancy only to those flows that absolutely require it!
  - It appears to be more practical to apply redundancy on Actuator as compared to Sensor flows
- End-to-end protection from talker to listener results in the highest integrity communication
- Dual-homing & temporal redundancy both offer very high robustness against soft errors where multi-homing additionally protects against a single hard error to/from the ECU
- Using IEEE 802.1CB as shown allows development migration from “no redundancy” to “best redundancy” with progressive steps (firmware updates) in between
  - This allows re-use of existing ECUs without needing to rewrite/redesign everything in one go

SECURE CONNECTIONS
FOR A SMARTER WORLD

Verifying analytical model using Simulation-Based Fault-Injection

- Simulation with RTaW-Pegase, Bit Error Rate = 10^{-7}

Solution	Packet	Loss rate by analysis	Loss rate by simulation
1	Largest	2.44E-3	2.45E-3
	Smallest	1.02E-4	1.04E-4
2	Largest	2.45E-3	2.47E-3
	Smallest	1.02E-4	1.02E-4
3	Largest	1.33E-5	1.34E-5
	Smallest	2.36E-8	2.14E-8
4	Largest	2.36E-5	2.35E-5
	Smallest	4.19E-8	4.14E-8
5	Largest	1.23E-3	1.23E-3
	Smallest	5.12E-5	5.12E-5

Simulation not suited for realistic BERs: e.g., one packet loss every 870 days on average, for a 1ms stream (largest packet size) with BER=10^{-10}