AUTOMOTIVE SYSTEM REQUIREMENTS ON TRAFFIC SHAPING.

KARL BUDWEISER, BMW GROUP
DR. JÖRN MIGGE, RTAW
TSN/A CONFERENCE 2024
2 OCTOBER, STUTTGART

AGENDA.

1. **Motivation for Traffic Shaping**
2. A Scene-based Analysis of Shaping Approaches
3. Limitations of TSN Shapers
4. Topology-Adequate Traffic Preconditioning with HTB
5. Summary & Prospects

BMW GROUP

RTaW RealTime-at-Work

TSN Technology & Applications

MOTIVATION FOR TRAFFIC SHAPING.
WHEN, WHERE, AND HOW?

Let’s lay the foundation for a thorough understanding of when, where and how to apply traffic shaping.

- When do bursts occur?
  - when (repeatedly) produced data does **not fit into one frame**
  - and these frames are **queued faster** than they can be **emitted**
  - then we get bursts of frames that are sent back-to-back and must be stored until transmission
- Why is shaping important? Non shaped higher priority bursty traffic …
  - leads to higher memory requirements inside the network, if frame losses must be avoided
  - makes lower priority traffic wait longer, which leads to further bursts and strong delay variations
- Can we always shape?
  - **To shape** a stream of frames **means to delay** frames, but **can we afford to delay them**?
  - Since **links must not be overloaded**, it must be possible to **emit the frames of every bursts in less than the burst-period**; this means that there is some **margin** for delaying
- Where should we shape? What happens if we do not shape bursty traffic (every where) in the system?

AGENDA.

1. Motivation for Traffic Shaping
2. **A Scene-based Analysis of Shaping Approaches**
3. Limitations of TSN Shapers
4. Topology-Adequate Traffic Preconditioning with HTB
5. Summary & Prospects

BMW GROUP

RTaW RealTime-at-Work

TSN Technology & Applications

WHAT HAPPENS IF WE DO NOT SHAPE PERIODIC BURST TRAFFIC?

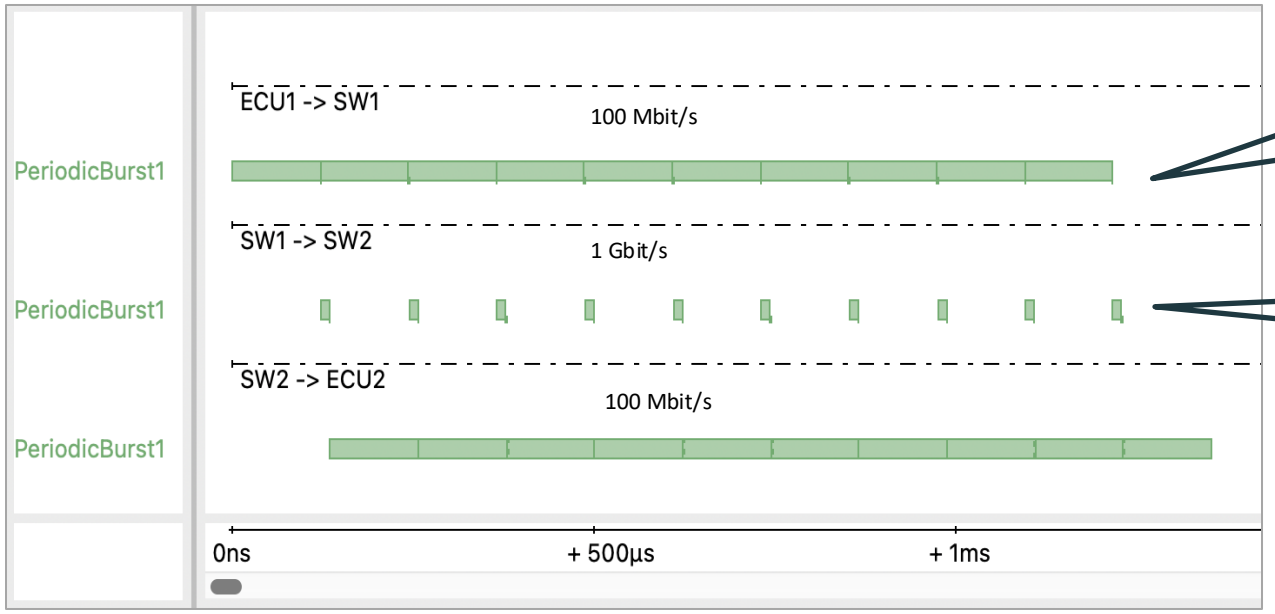

1: Frames of the burst are **queued** much **faster than** they can be **transmitted**

2: Frames are **transmitted faster than** they are **queued**

⇒ 100 Mbit/s **shaping** of first link is visible on following 1000 Mbit/s link

Source Node	Memory	Nbr Frames	Tx Rate	Load	Target Node
ECU1	10000 byte	10 frames	100 Mbit/s	16,32%	SW1
SW1	1000 byte	1 frames	1000 Mbit/s	1,63%	SW2
SW2	2000 byte	2 frames	100 Mbit/s	16,32%	ECU2

Some natural shaping occurs, but alone it has disadvantages …

A SCENE-BASED ANALYSIS OF DIFFERENT APPROACHES FOR TRAFFIC SHAPING.
LINK SPEED REDUCTION

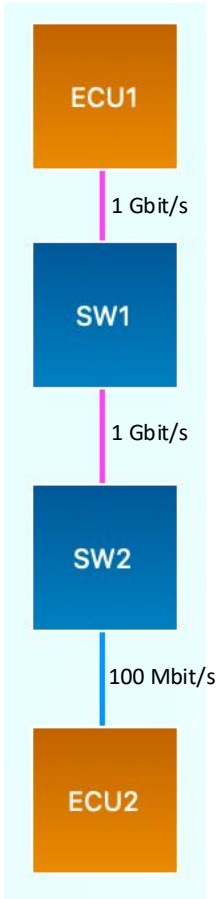

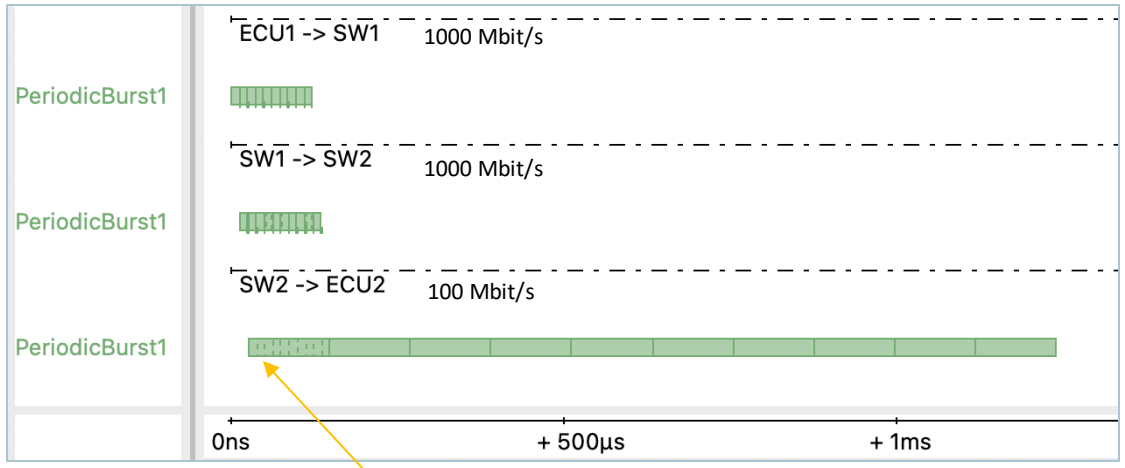

without shaping: burst of frames needs to be stored again

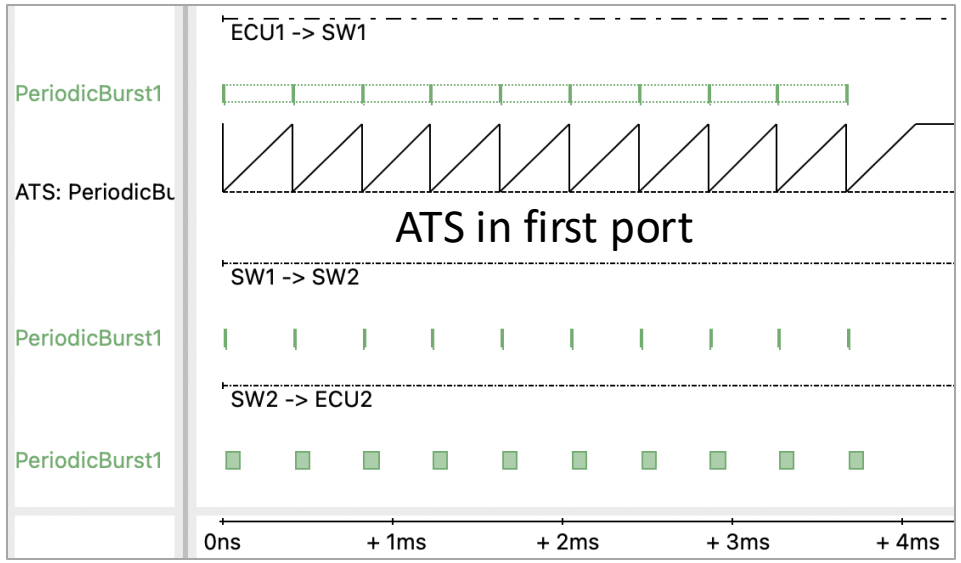

with (ATS) shaping: **burst of frames stored only once** (in the port of the ECU)

	Non shaped		Shaped				
Source Node	**Memory**	**Nbr Frames**	**Memory**	**Nbr Frames**	**Tx Rate**	**Load**	**Target Node**
ECU1	10000 byte	10 frames	10000 byte	10 frame	1000 Mbit/s	SW1	1,63%
SW1	2000 byte	2 frames	1000 byte	1 frame	1000 Mbit/s	SW2	1,63%
SW2	10000 byte	10 frames	1000 byte	1 frame	100 Mbit/s	ECU2	16,32%

Note: CBS also allows to obtain a similar effect.

A SCENE-BASED ANALYSIS OF DIFFERENT APPROACHES FOR TRAFFIC SHAPING.
JOINING STREAMS (WITHOUT LINK SPEED CHANGE)

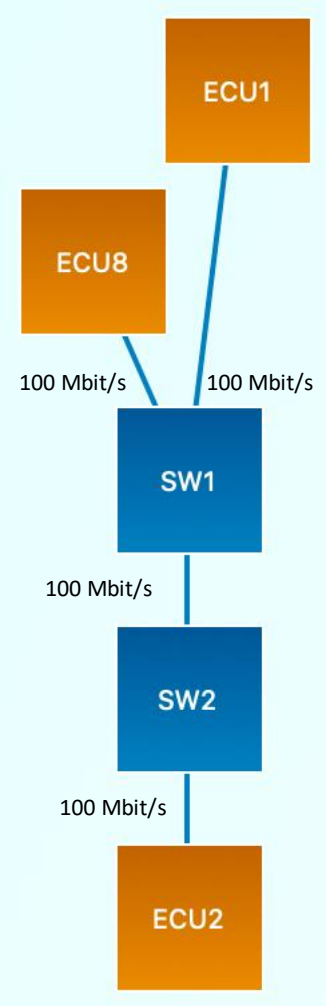

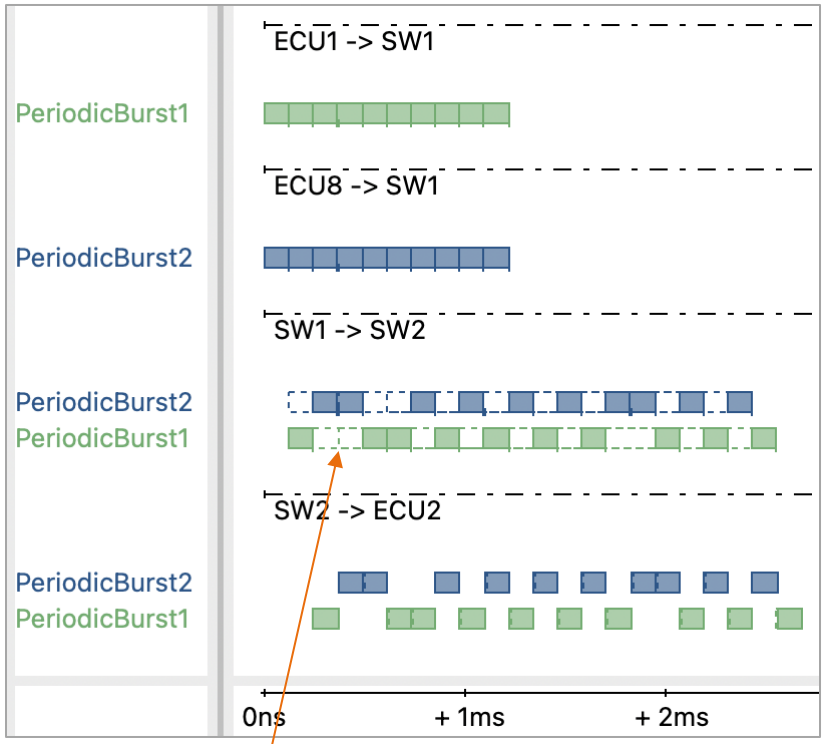

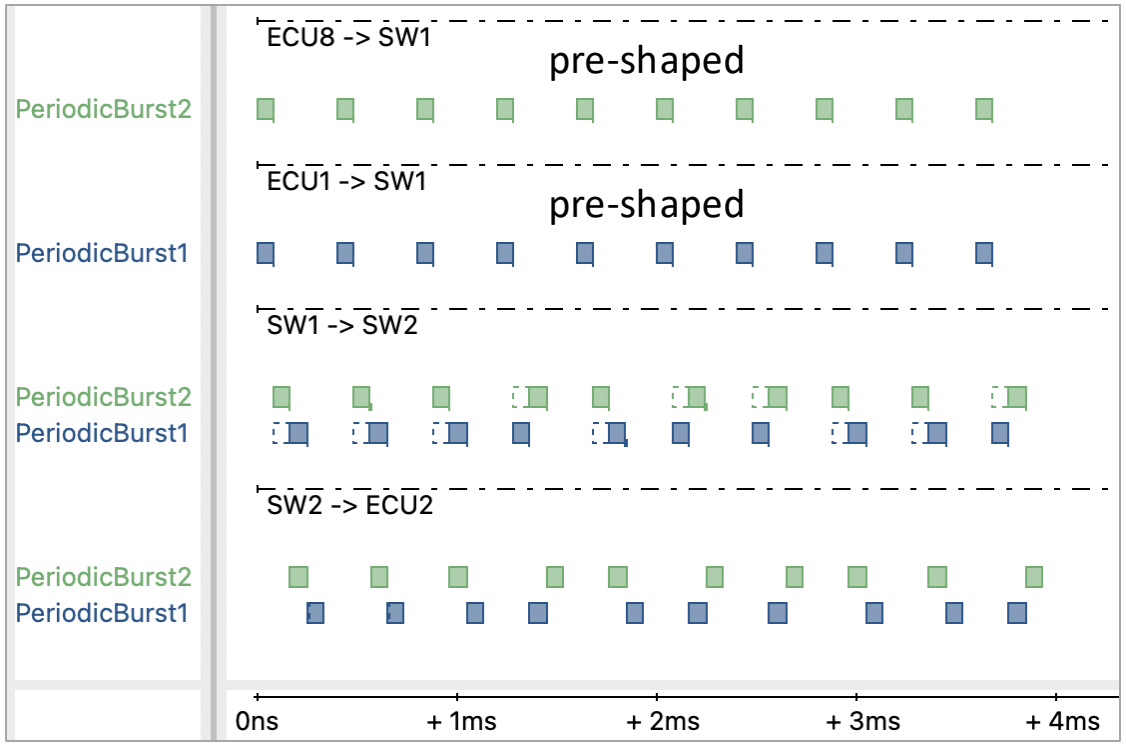

without shaping: burst of frames need to be stored again

	Non shaped		Shaped				
Source Node	Memory	Nbr Frames	Memory	Nbr Frames	Tx Rate	Load	Target Node
ECU1	10000 byte	10 frames	10000 byte	*10 frame	100 Mbit/s	16,32%	SW1
ECU8	10000 byte	10 frames	10000 byte	*10 frame	100 Mbit/s	16,32%	SW1
SW1	12000 byte	12 frames	2000 byte	2 frames	100 Mbit/s	32,64%	SW2
SW2	2000 byte	2 frames	2000 byte	2 frames	100 Mbit/s	32,64%	ECU2

*: frames are stored in the ECU and not in the port

EFFECTS ON LOWER PRIORITY TRAFFIC (1)

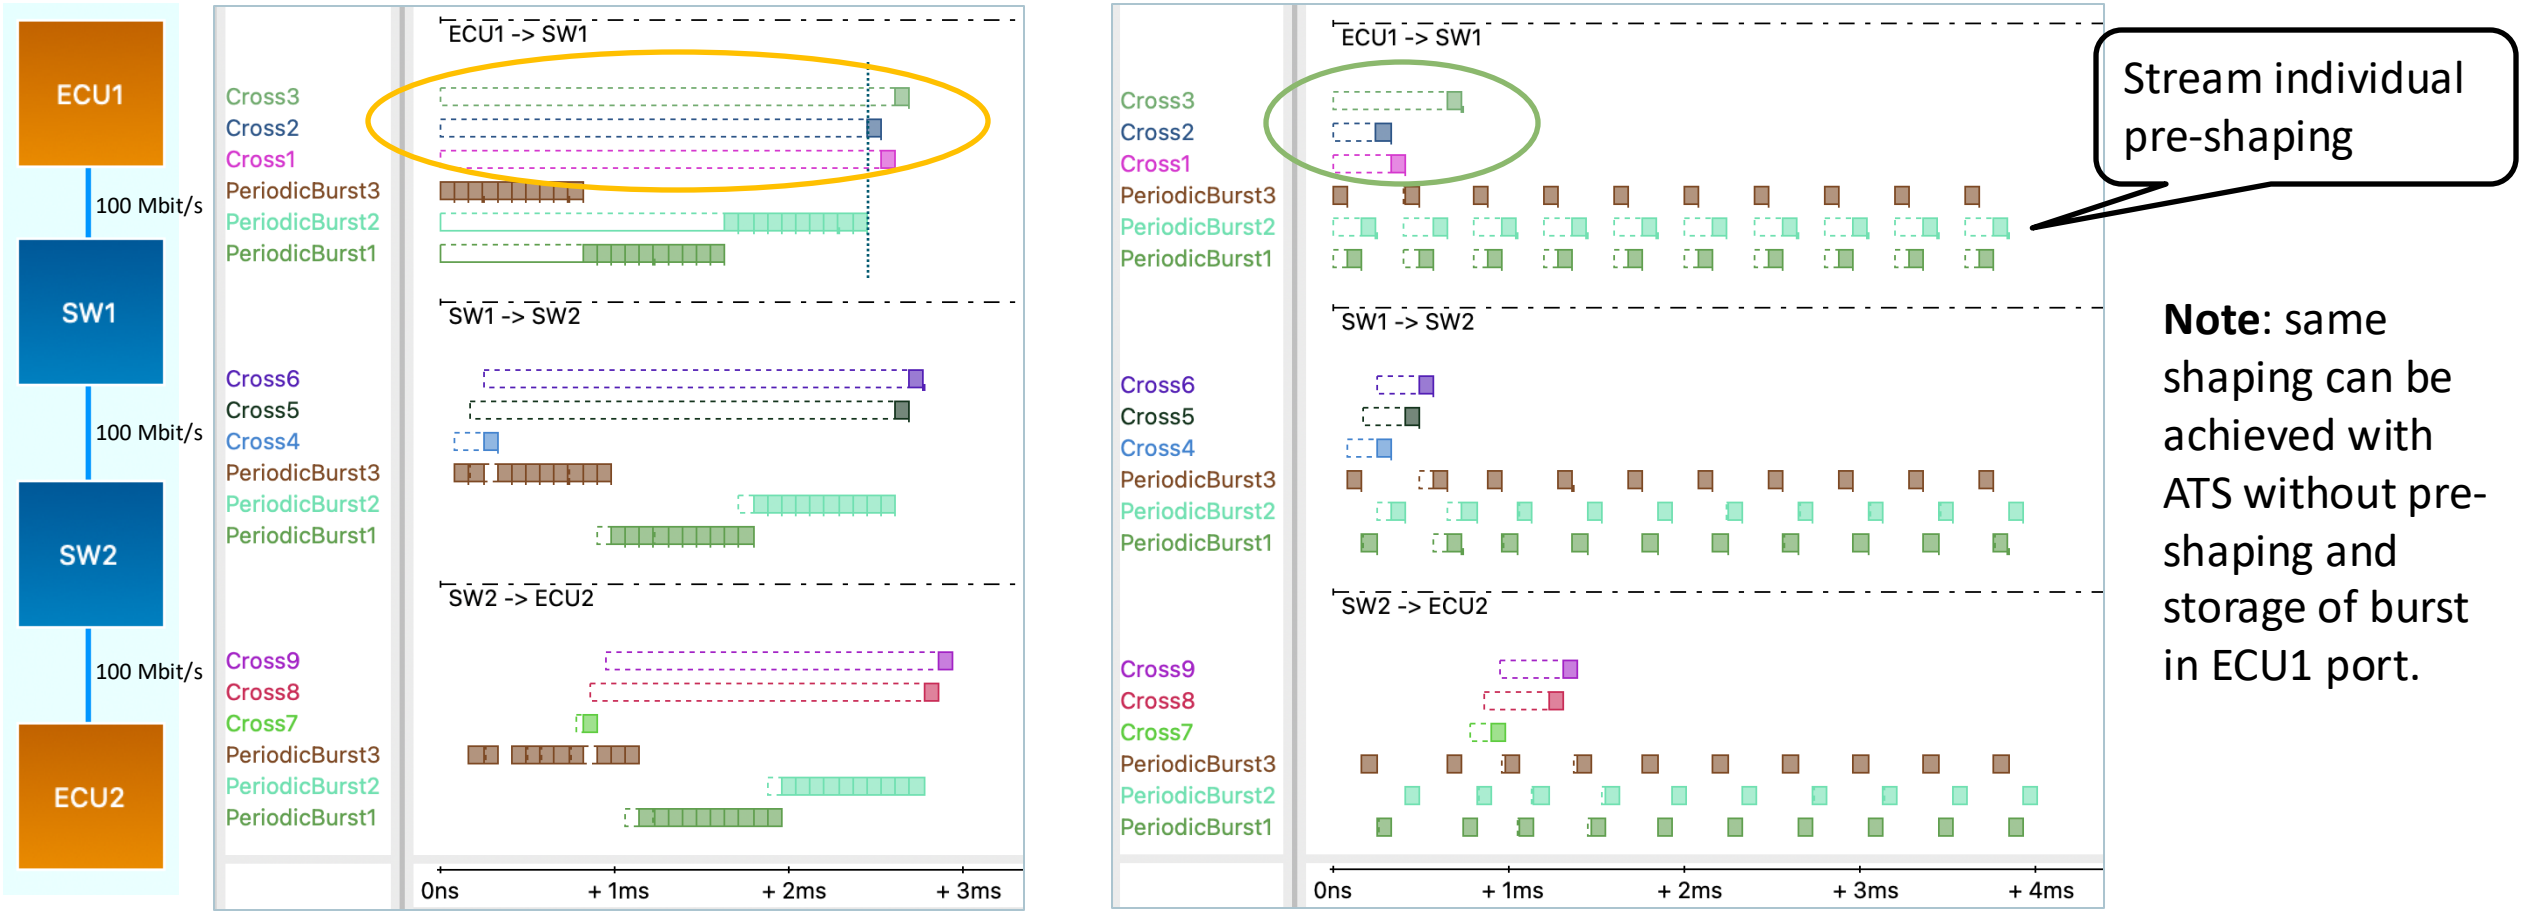

Note: same shaping can be achieved with ATS without pre-shaping and storage of burst in ECU1 port.

(pre-) **shaped higher priority periodic bursts** *may* **lower worst-case delays** for **lower priority frames**

EFFECTS ON LOWER PRIORITY TRAFFIC (2)

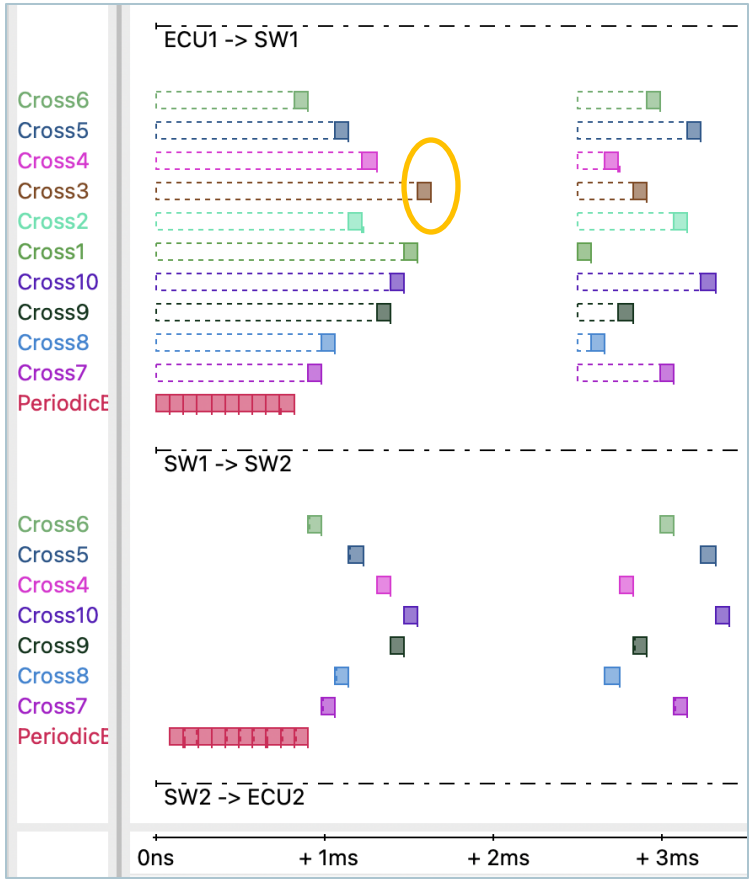

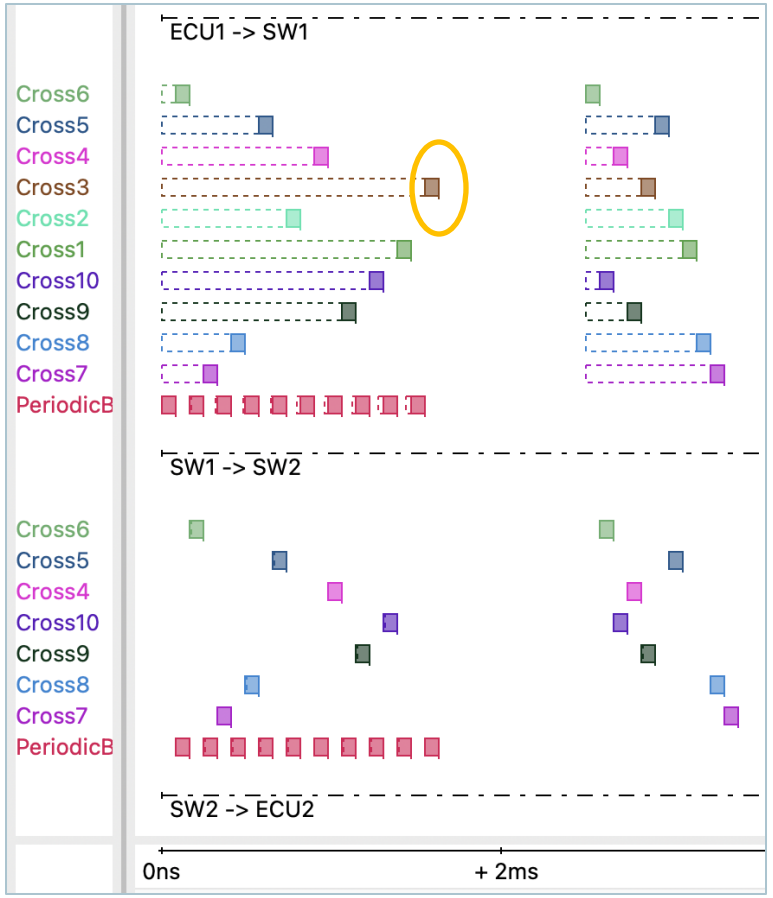

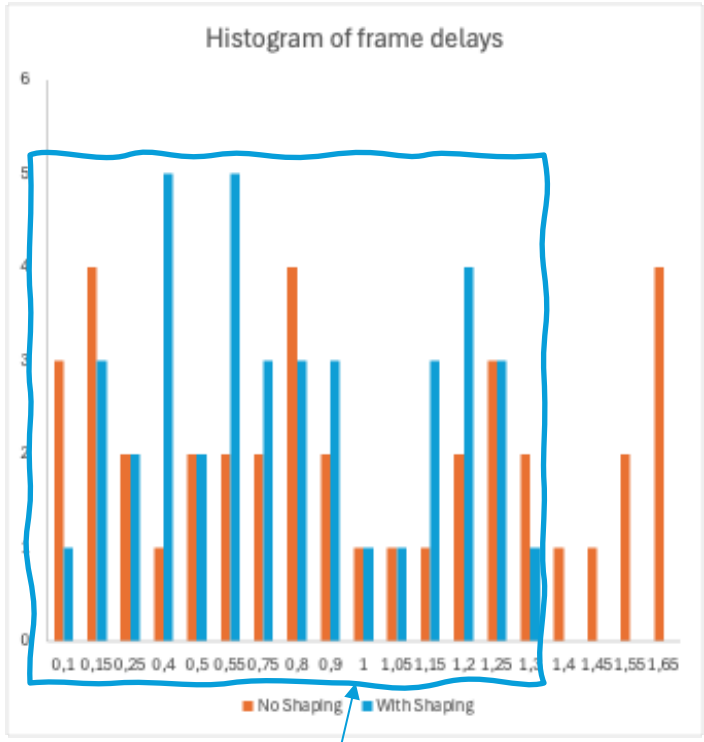

However, “With shaping” of higher priority periodic bursts:
⇒ **high delays less frequent** for **lower priority frames**

But if we have many lower priority frames and short higher priority periodic burst: **worst-case delays might not become shorter** ….

PRE- SHAPING + FIRST PORT SHAPING

ECU port

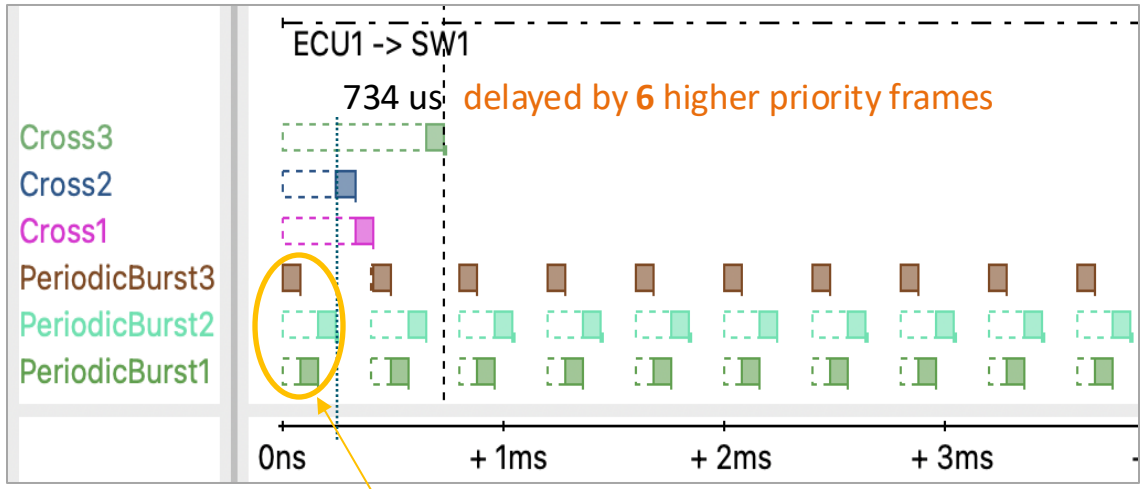

Multiple, independently (pre-)shaped streams can produce micro-bursts

ECU port

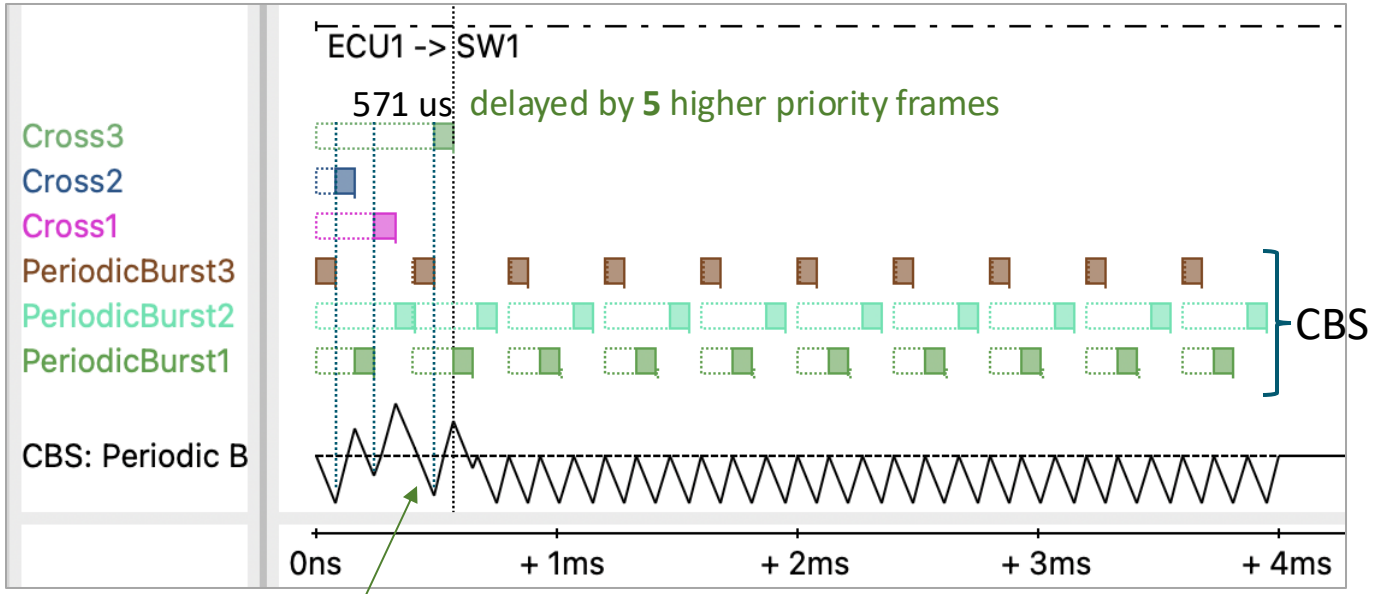

CBS class shaping may produce more/earlier transmit opportunities for lower priority frames

Note: this cannot be achieved with ATS because network ingress shapers work per stream.

FIRST PORT VS END-TO-END SHAPING

RTaW RealTime-at-Work

Switch port

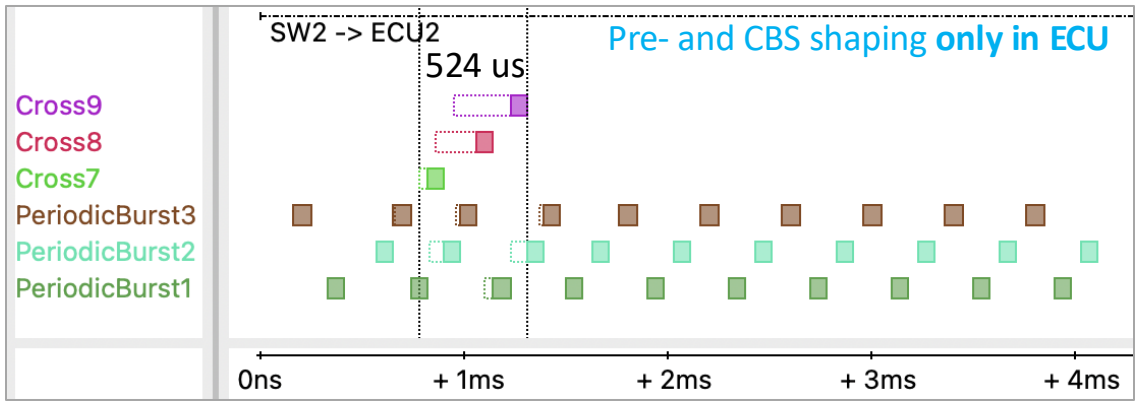

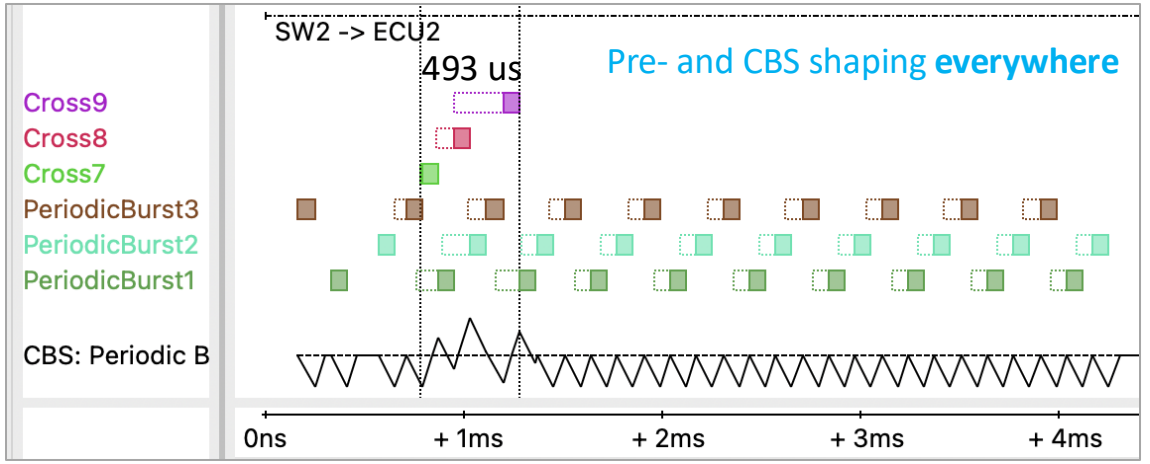

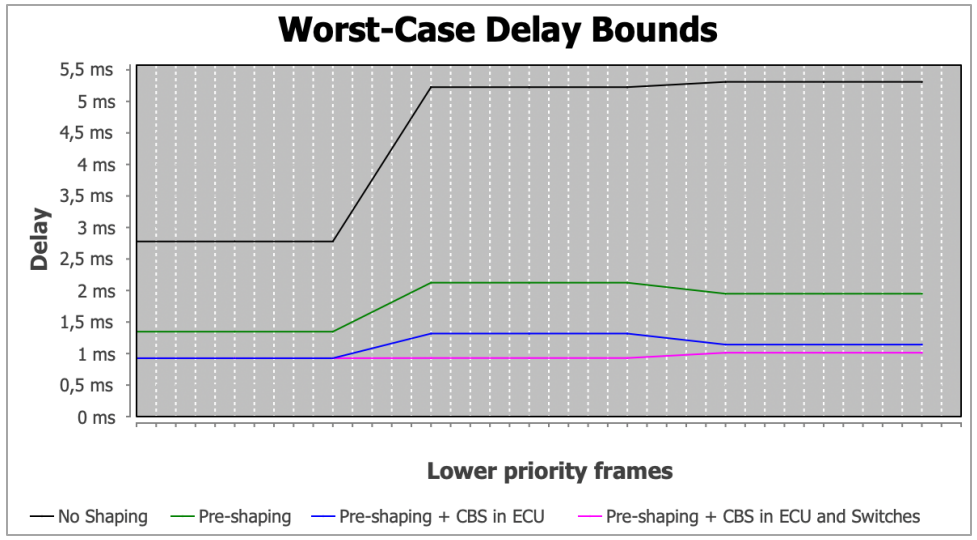

End-To-End shaping:

- may improve delays
- but may require more port memory

⇒ trade-off to find

No general answer: must be studied with the help of **worst-case analysis tools**

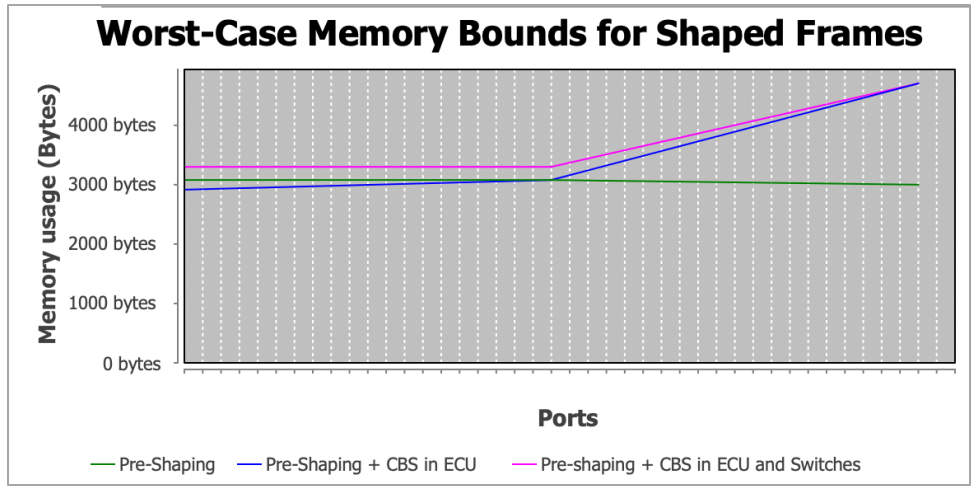

AGENDA.

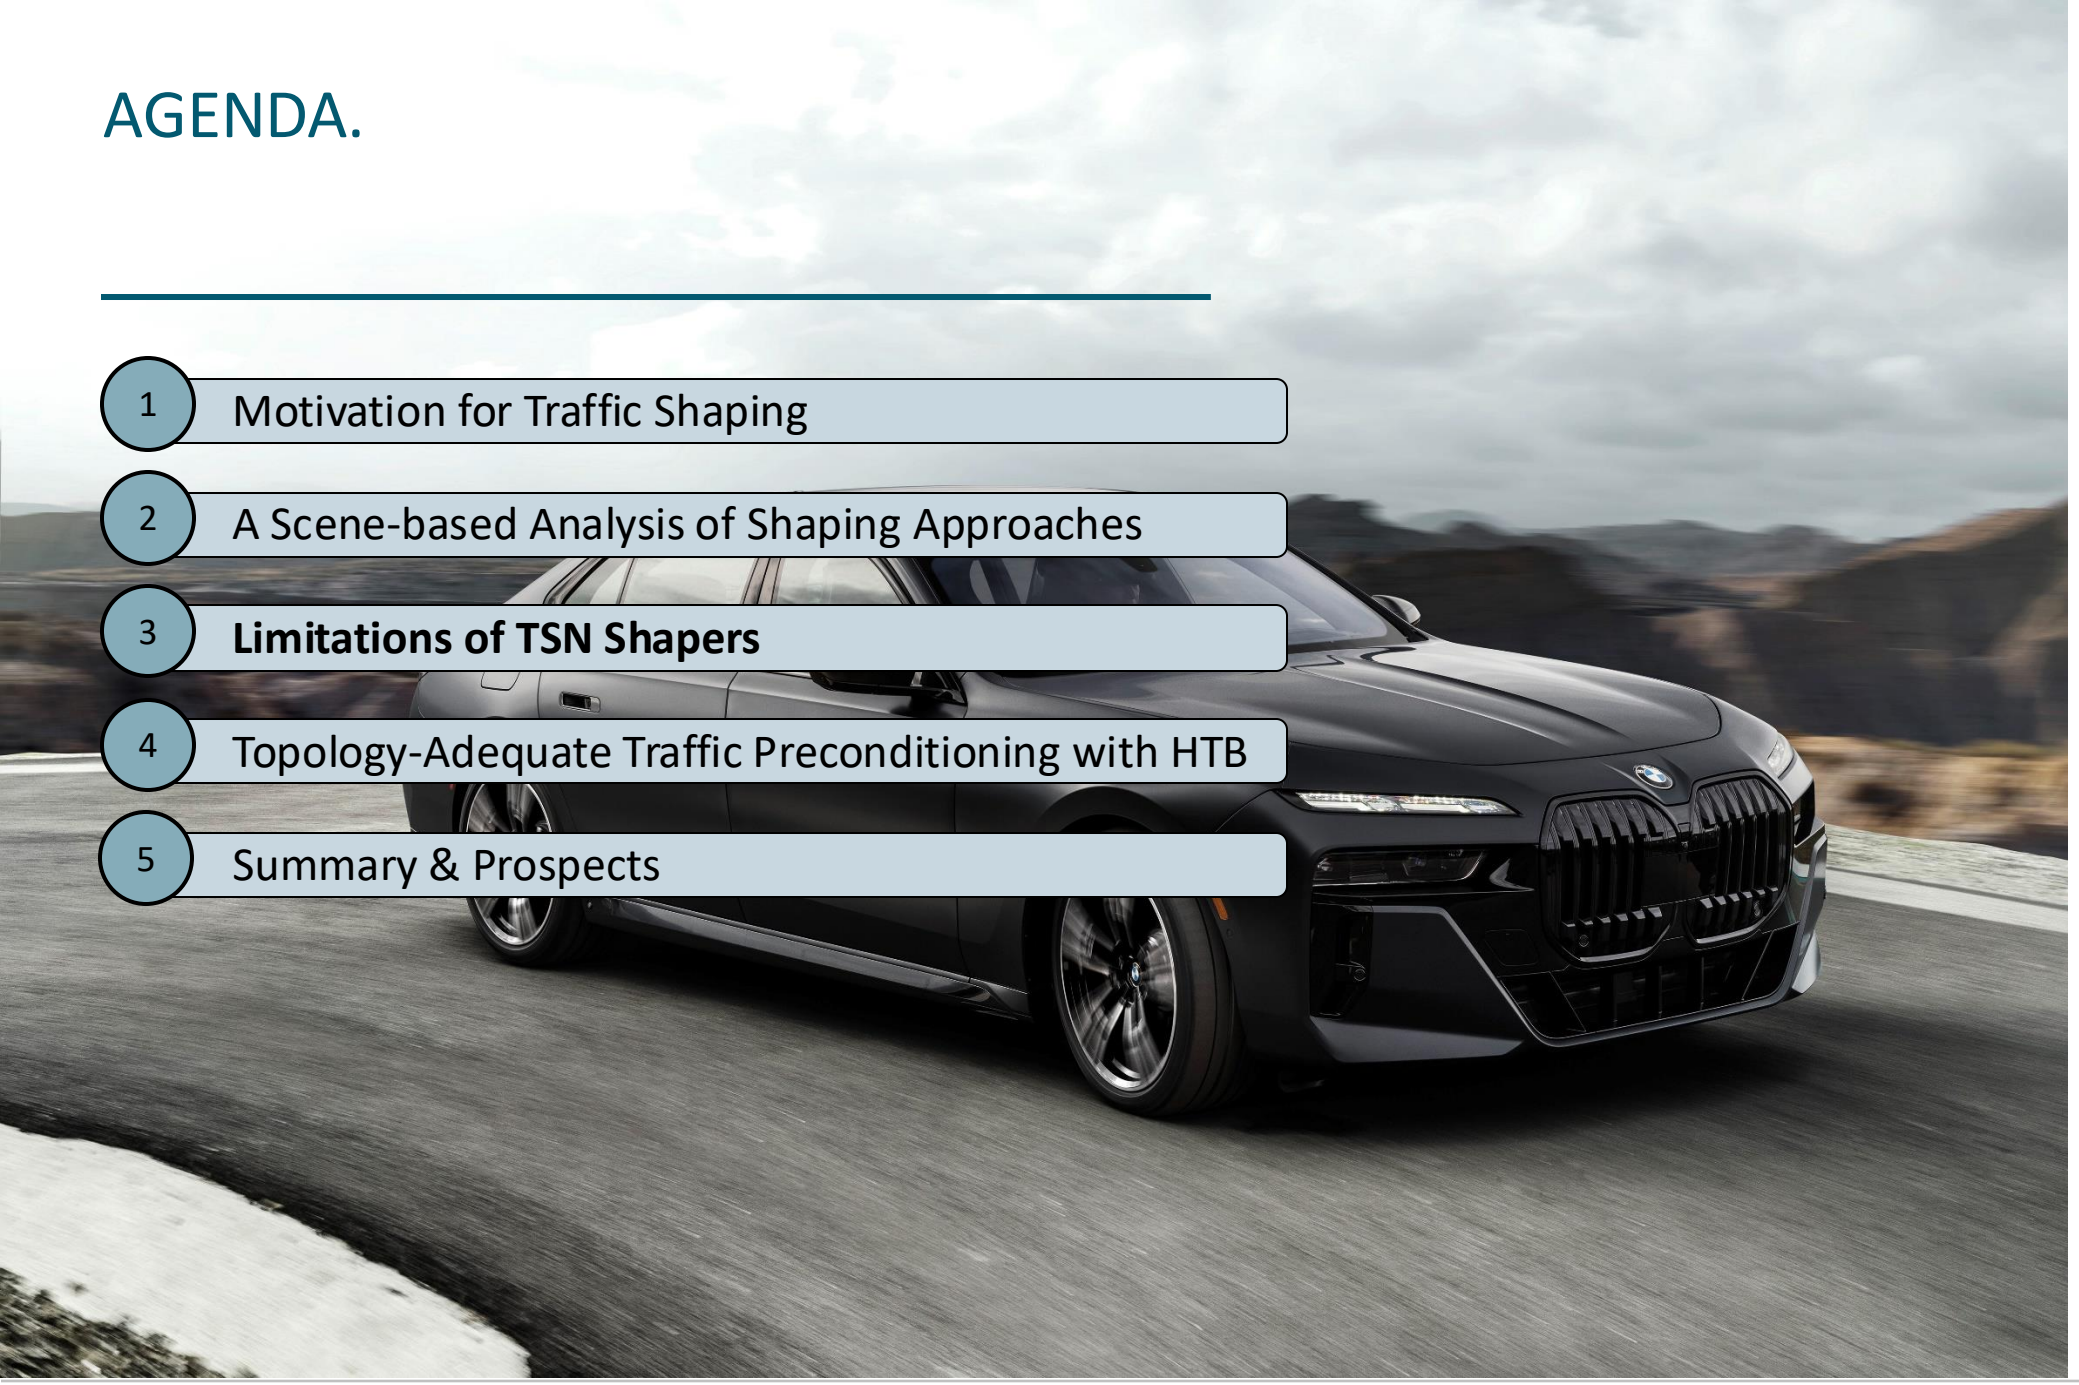

1. Motivation for Traffic Shaping
2. A Scene-based Analysis of Shaping Approaches
3. **Limitations of TSN Shapers**
4. Topology-Adequate Traffic Preconditioning with HTB
5. Summary & Prospects

BMW GROUP

RTaW RealTime-at-Work

TSN Technology & Applications

SYSTEM-LEVEL CONSIDERATIONS AND SOLUTIONS.

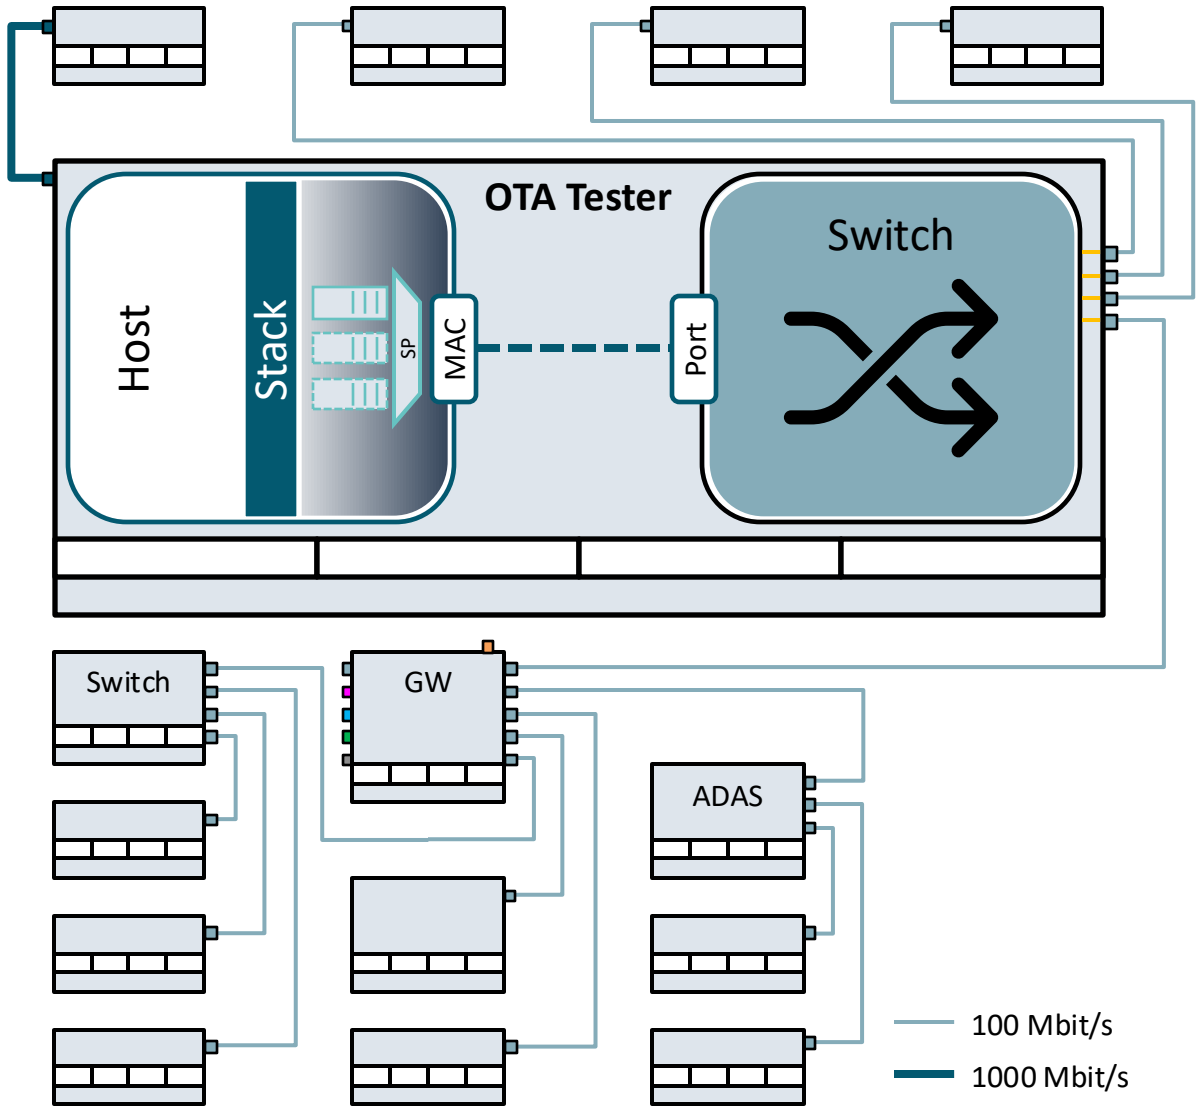

Use Case and Requirements:

- For OTA update the **Tester** streams flash data via dedicated TCP connections to all diagnosable ECUs
- To reduce programming time all links **shall** run at the highest possible rate w/o overloading any link aggregating streams

Problems:

- Mismatching speed-grades may introduce **bottlenecks**:
  → network congestions & packet loss

Solution Option #1:

- Use a single **shaper instance** to throttle down all diagnostic communication to the lowest speedgrade used in the system
  + Simplistic shaper configuration
  - "Over-shaping" that tends to cause inefficiency by wasting bandwidth

SYSTEM-LEVEL MEASURES FOR PERFORMANCE TUNING.

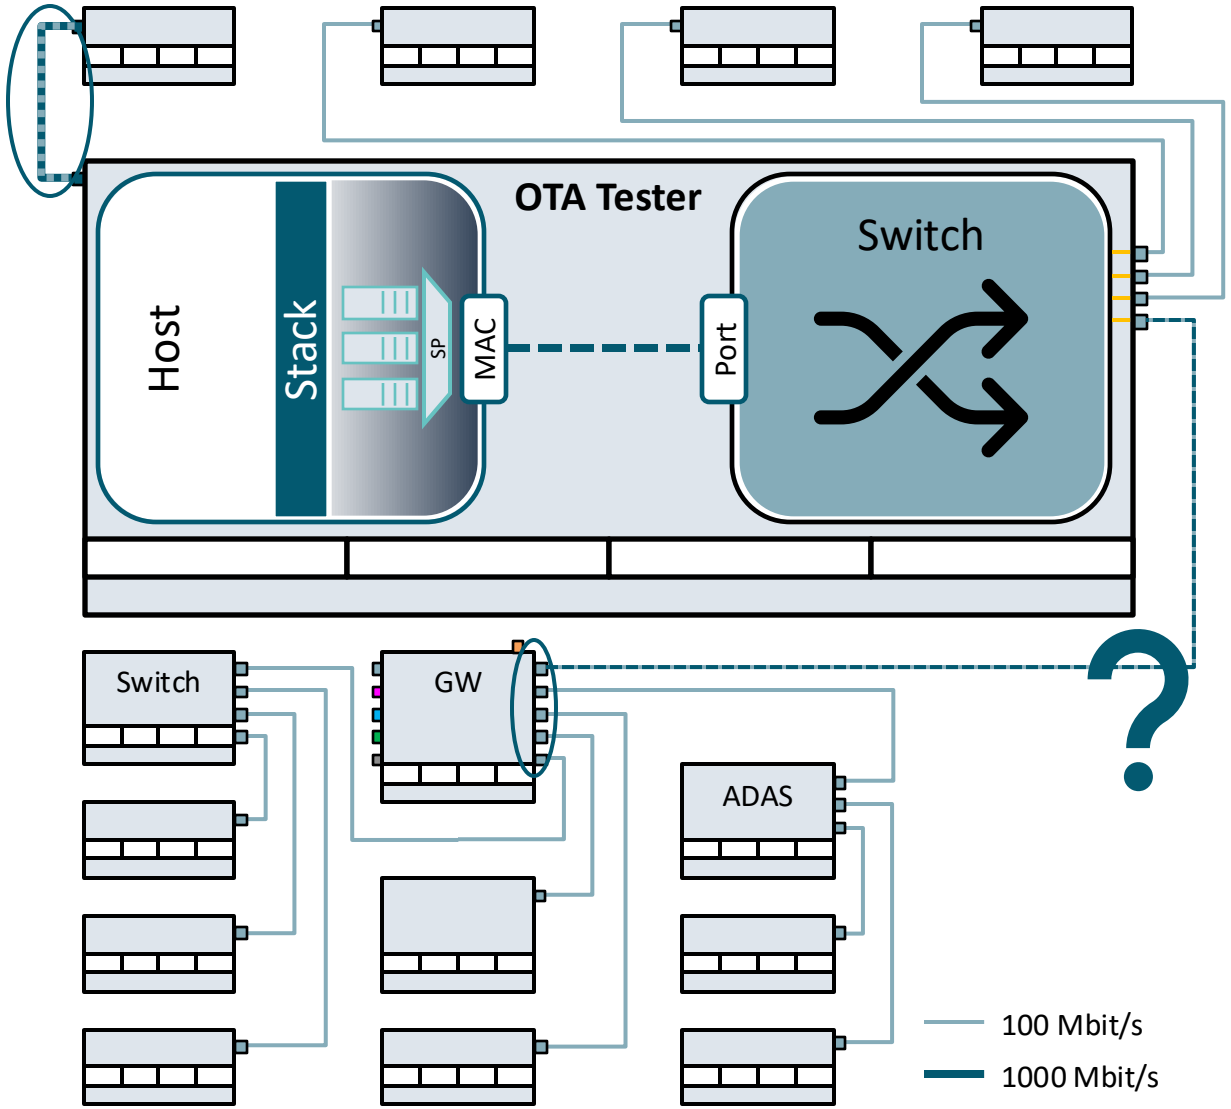

Use Case and Requirements:

- For vehicle programming the **OTA Tester** streams flash data via dedicated TCP connections to all diagnosable ECUs
- To reduce programming time all links **shall** run at the highest possible rate w/o overloading any link aggregating streams

Problems:

- Mismatching speed-grades may introduce **bottlenecks**:
  → network congestions & packet loss

Solution Option #2:

- Use a single **shaper instance for each switch port** to throttle down diagnostic communication to the lowest speed-grade used on the respective branch
  - \+ Simplistic, but more **topology-dependent** configuration
  - \- Not always feasible due to HW-constraints (#shapers)
  - \- Unfair due to starvation of low-prior traffic classes
  - \- better but risk of no optimal bandwidth utilization

AGENDA.

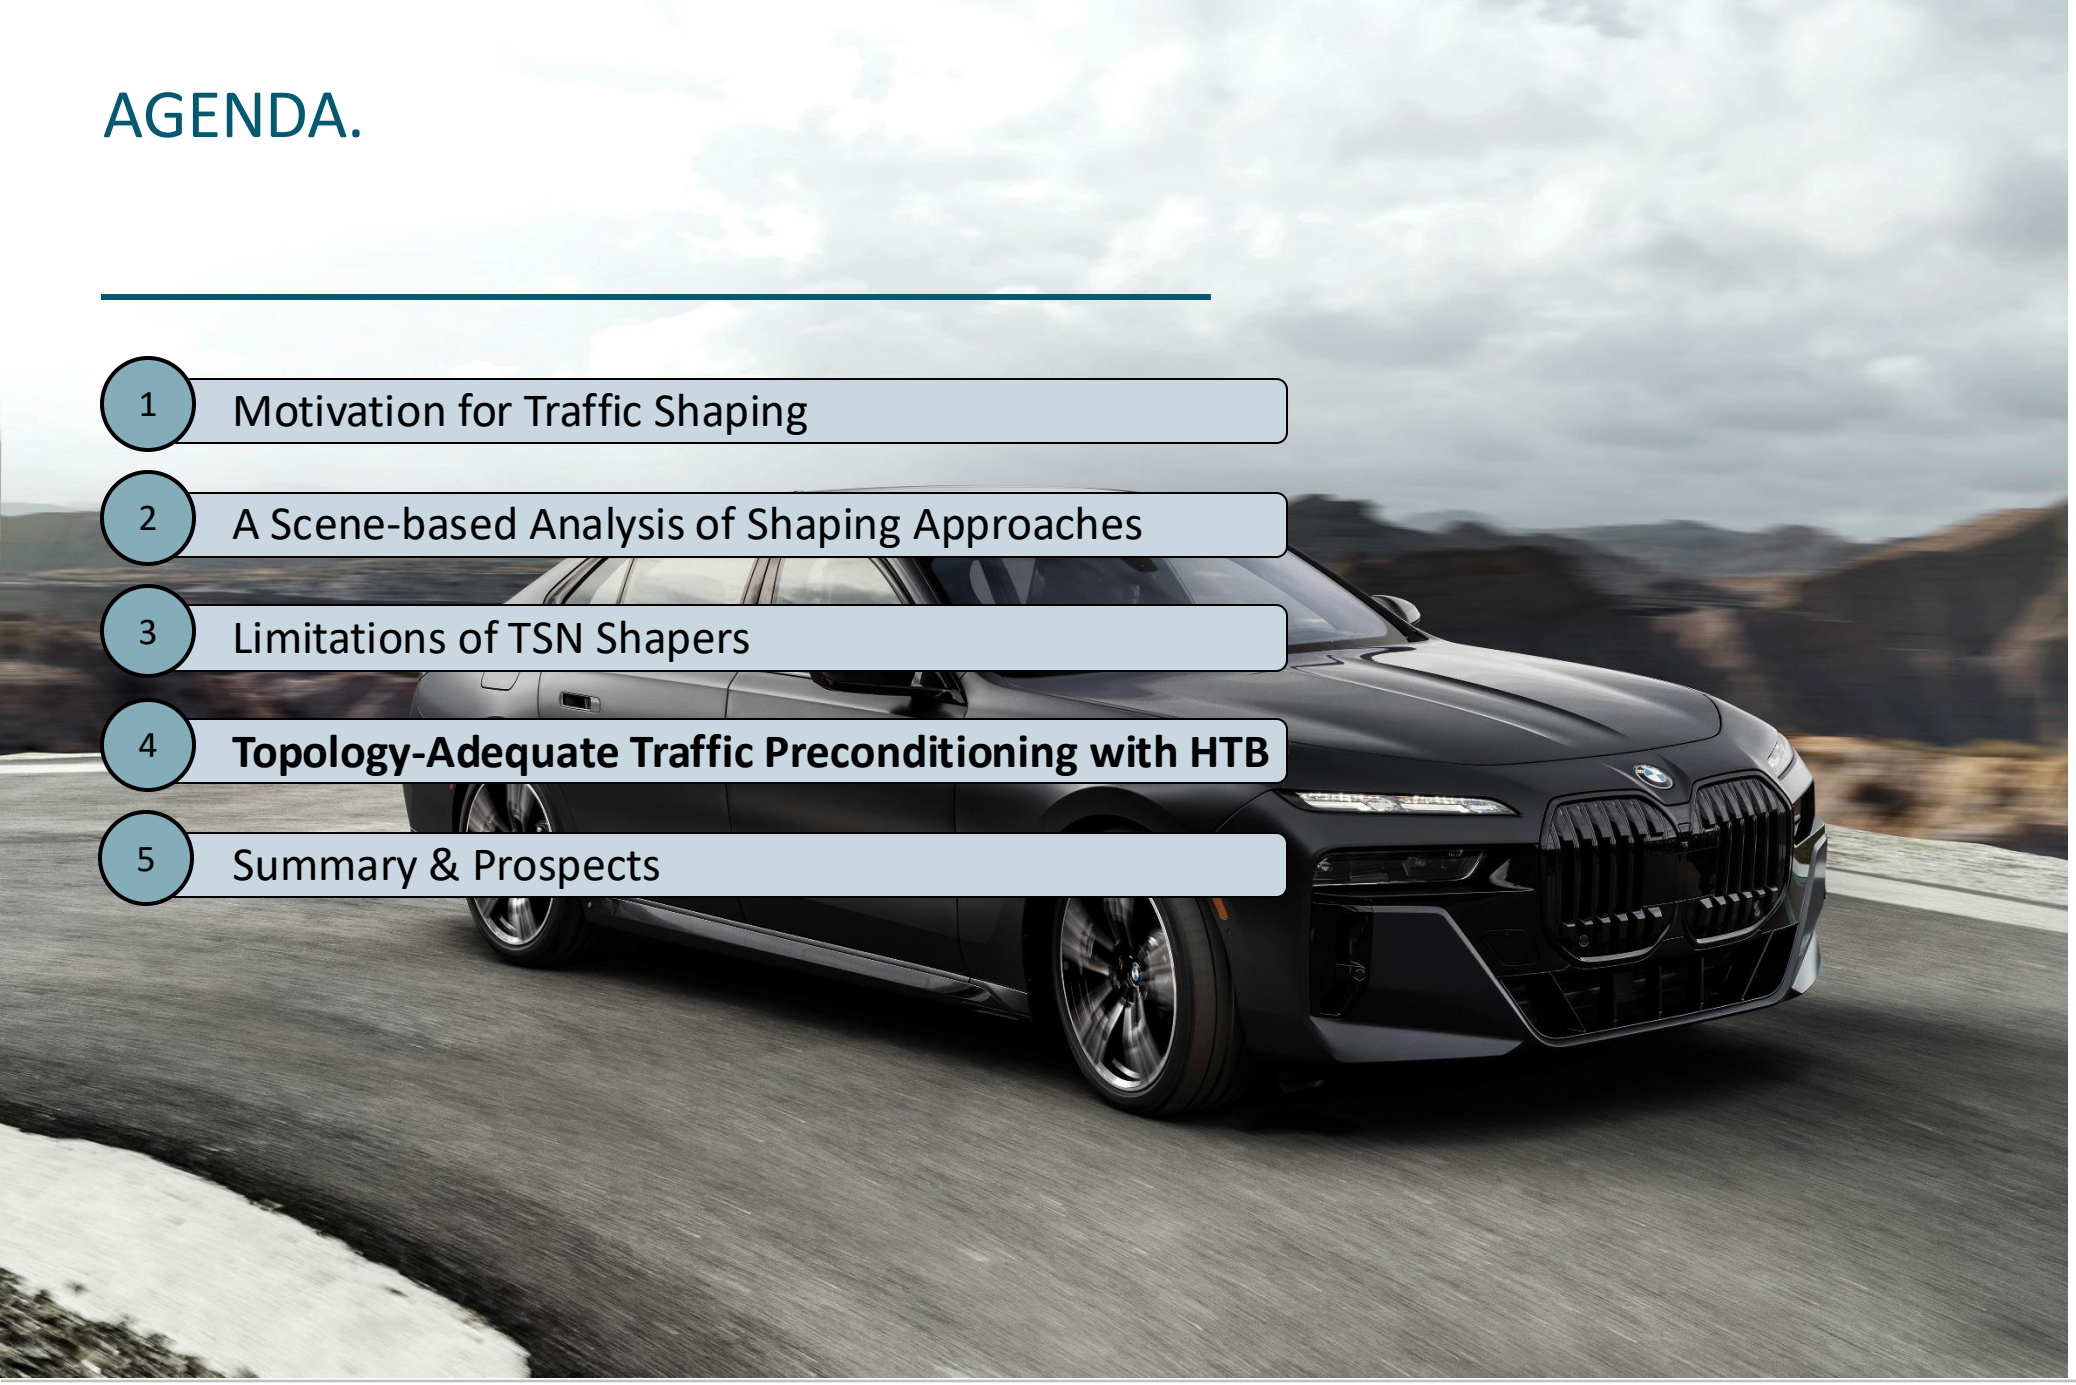

1. Motivation for Traffic Shaping
2. A Scene-based Analysis of Shaping Approaches
3. Limitations of TSN Shapers
4. **Topology-Adequate Traffic Preconditioning with HTB**
5. Summary & Prospects

BMW GROUP

RTaW RealTime-at-Work

TSN Technology & Applications

THE OPERATING PRINCIPLE OF THE HTB.

HTB is a **class-based shaper**, where child classes can be uniquely identified by the their class-name:

```
root:strcat(parent,number)
```

Operating Principle and Constraints:

- Child classes can be added to (parent) classes
- Traffic can be mapped to every class, but preferably to classes w/o childs (‘leafs’)
- Each class has the rates `(r,ceil)`, where `r` serves as a **bandwidth-reservation** and `ceil` as a **rate-limitation**.
- Childs can only “borrow” bandwidth (BW) from their parents
- In order to meet all reservation, the following **constraint** must be met for any parent class
  - $\sum_{childs} r_{child} < r_{parent}$
  - $\sum_{childs} r_{child} < ceil_{parent}$

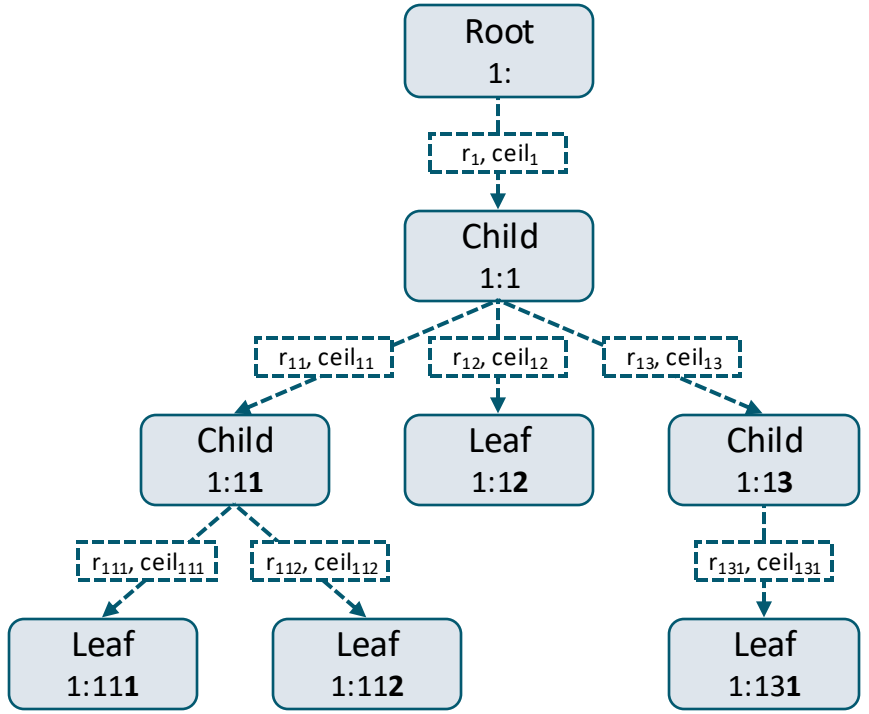

HTB can grant a BW-budget and rate-limit for any class

HTB can be used to grant a minimum bandwidth, while limiting the data rate of any link w/o overloading it

TOPOLOGY-ADEQUATE TRAFFIC PRECONDITIONING WITH HTB: SET-UP

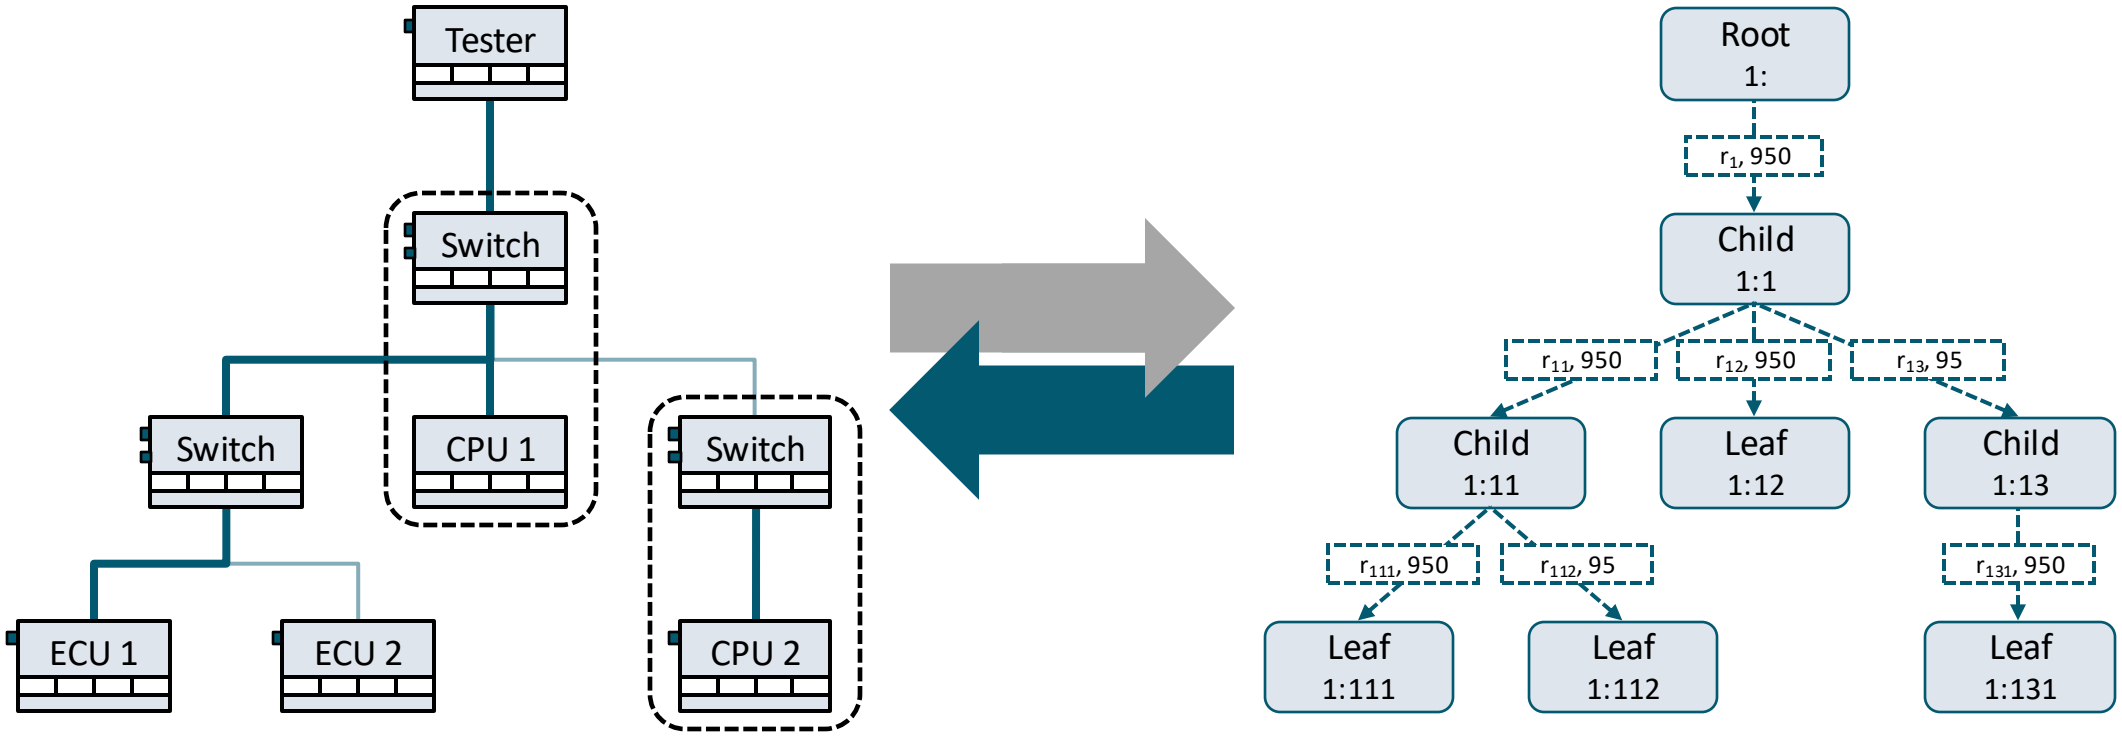

Links may operate at different speed-grades

Every link is rate-limited to ‘ceil’ [Mbps]

If the sender’s HTB configuration reflects the network topology with the sending ECU as root and every ceil-rate is set to the speed-grade of the corresponding link, links can be used to their full capacity w/o overloading any link

TIMING-BEHAVIOR OF TOPOLOGY-ADEQUATE TRAFFIC PRECONDITIONING.

Another Example:

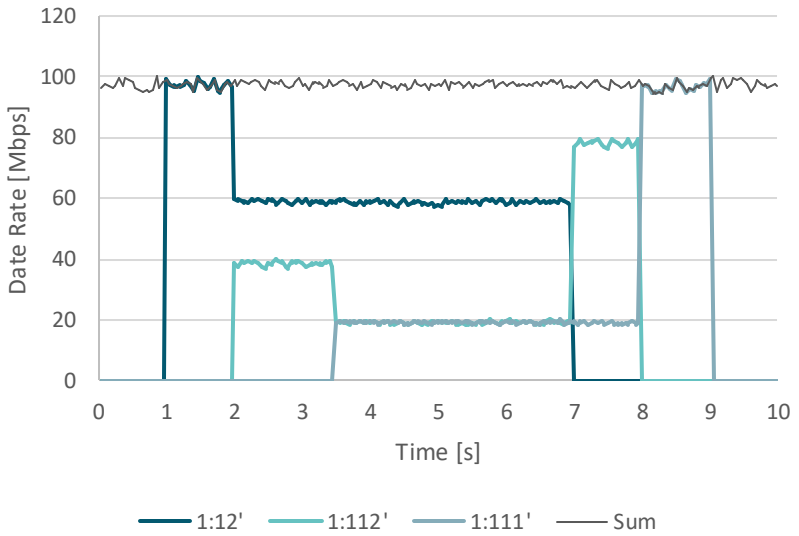

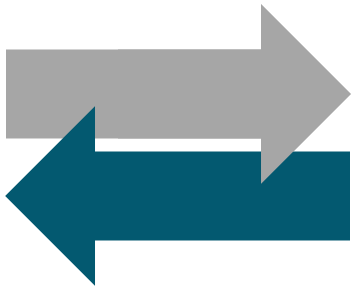

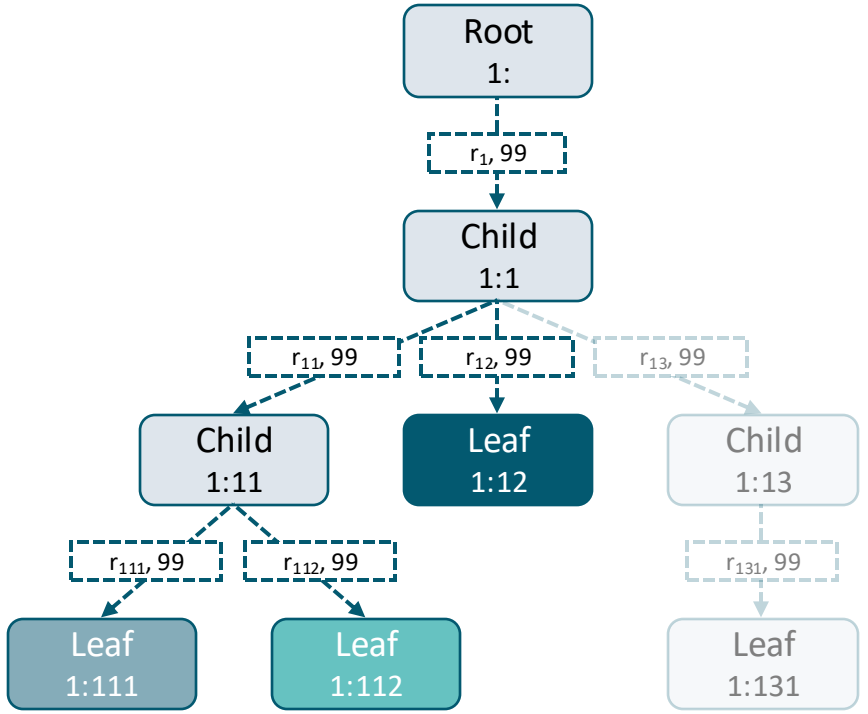

Rate ‘r’ and ‘ceil’ w/ prioritization of ‘leaf’

Asumptions:

- $r_1 >> r_{11} = 40$, $r_{12} = 60$, $r_{111} = r_{112} = 20$
- $prio_{11} \leq prio_{12}$, $prio_{111} < prio_{112}$

For each and every link a **bandwidth budget** can be specified by configuring the rate r. If no rates are specified the available bandwidth can either be shared **proportionally** (to ceil) or according to **priority**.

AGENDA.

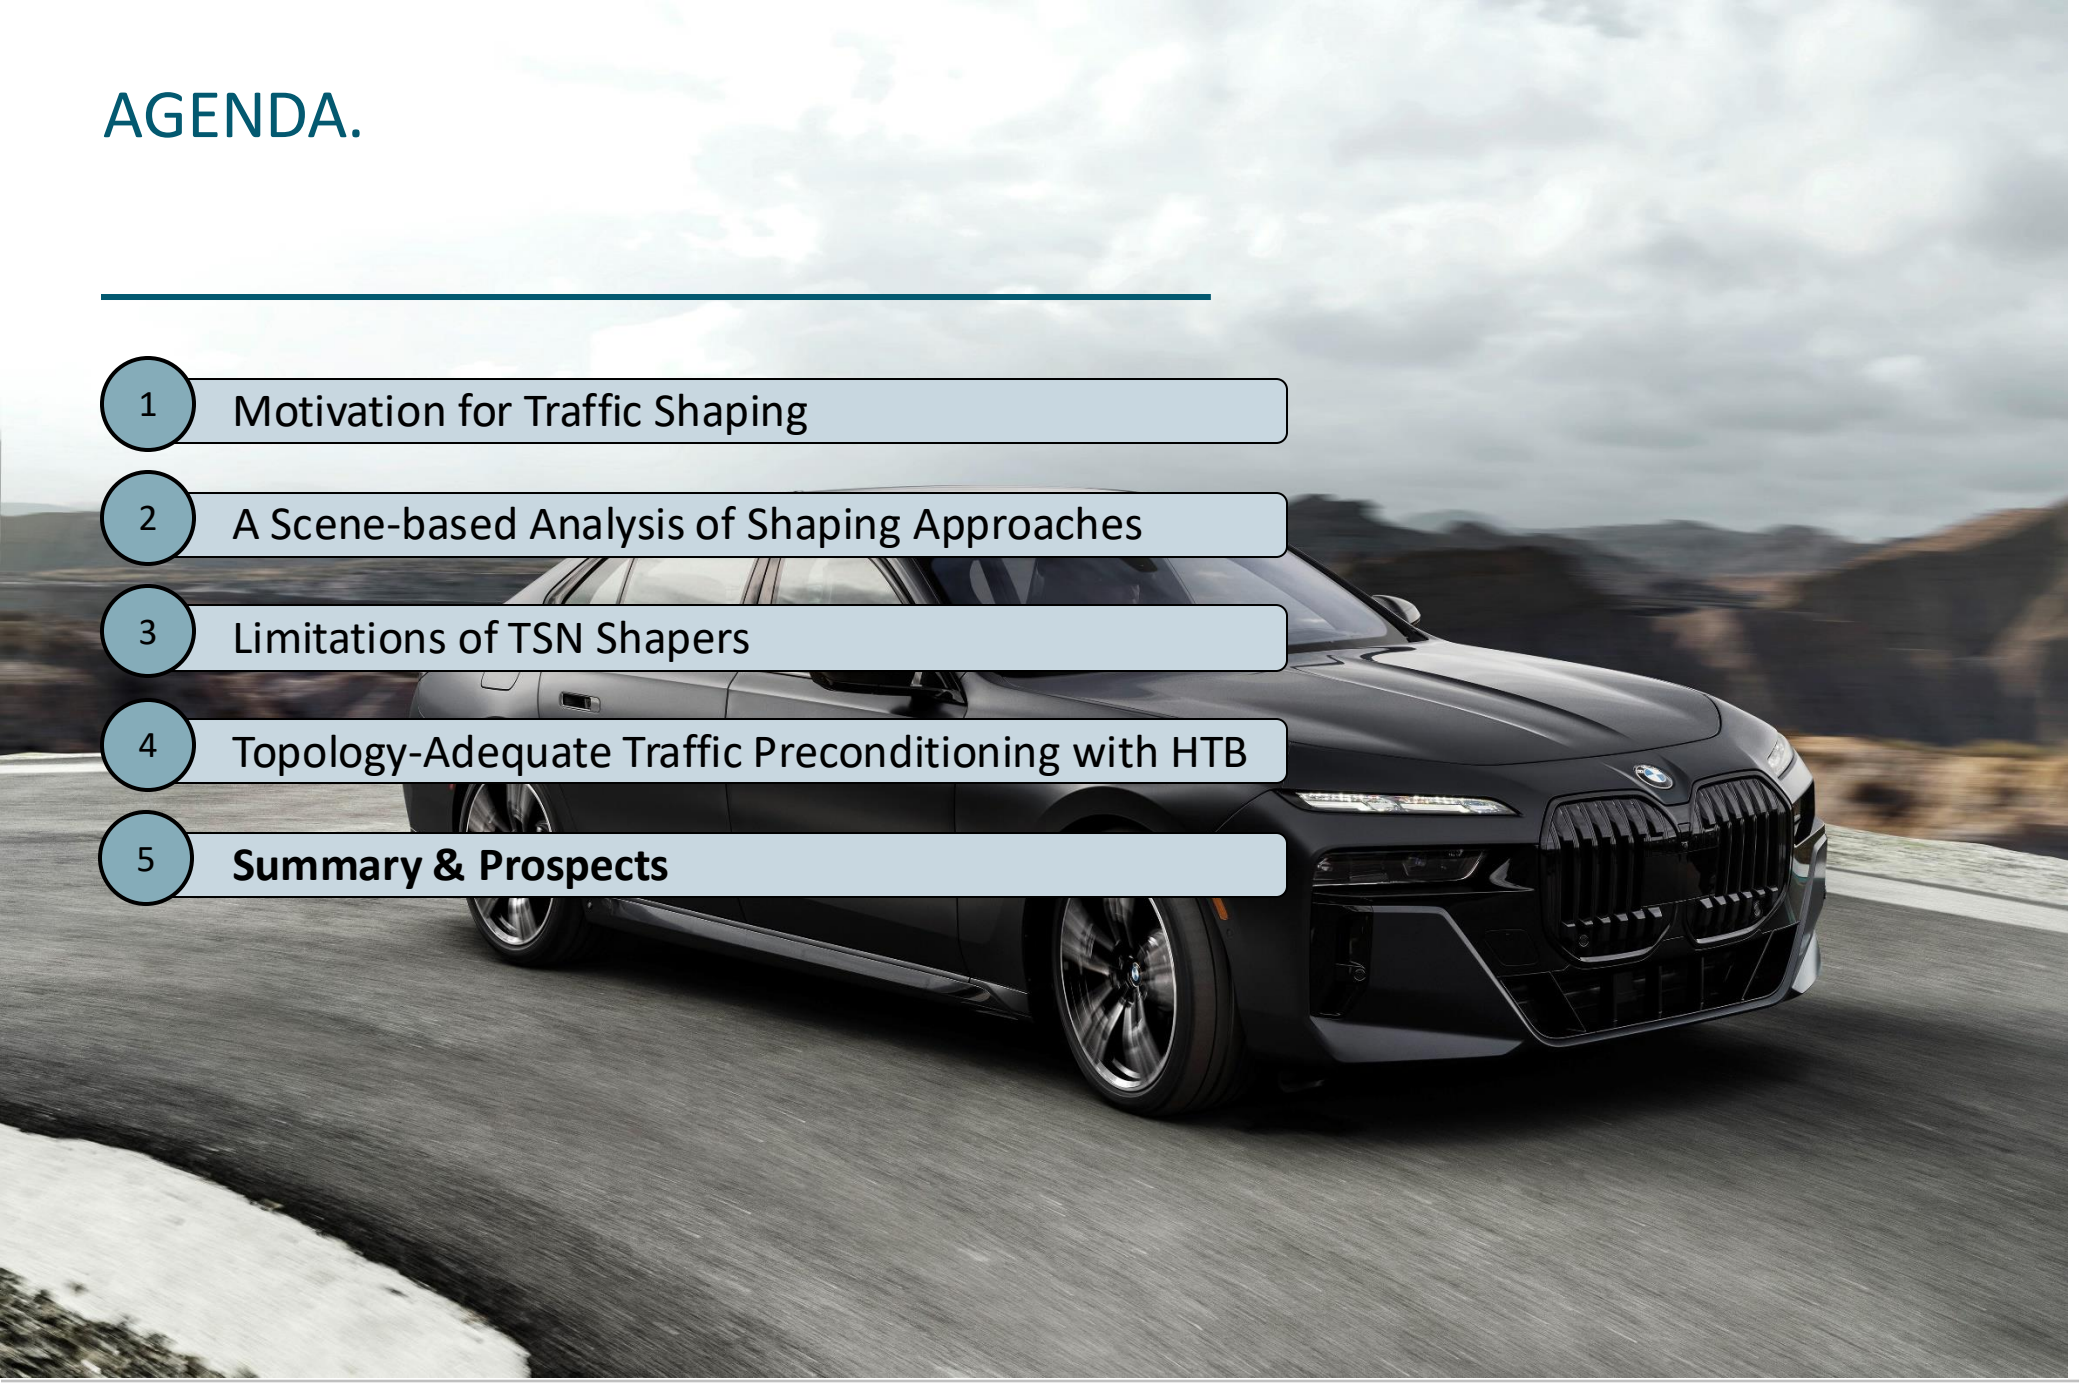

1. Motivation for Traffic Shaping
2. A Scene-based Analysis of Shaping Approaches
3. Limitations of TSN Shapers
4. Topology-Adequate Traffic Preconditioning with HTB
5. **Summary & Prospects**

BMW GROUP

RTaW RealTime-at-Work

TSN Technology & Applications

SUMMARY & PROSPECTS.

- With **Traffic Shaping**
  - **bursts** need to be **stored only once** in the system (reducing buffer costs especially for switches),
  - **lower latencies** (variations) can be achieved **for low priority traffic** due to regular transmission opportunities, though
  - shaped (bursty) high priority frames have **larger latencies but with lower variability**
- **E2E-Shaping** can improve timing-behavior to some extend, but requires more **configuration complexity** and **memory**.
- There is **no one-fits-all solution**, for different scenarios may require different approaches or configurations.
- Several very common **use cases** like
  - flash programming, or
  - Internet streaming

  require more **sophisticated traffic shaping** at the sending-side to achieve optimal performance.
- **HTB** provides a **hierarchical** and **holistic approach** that can be used to resolve issues classic **TSN shapers** cannot tackle.
- The HTB-configuration for **Topology-Adequate Traffic Preconditioning** is not only intuitive but also resource-optimal.

→ Let’s shape the future of in-vehicle networking together!


BMW
GROUP
MINI
RTaW
RealTime-at-Work
THANK YOU FOR YOUR ATTENTION.