



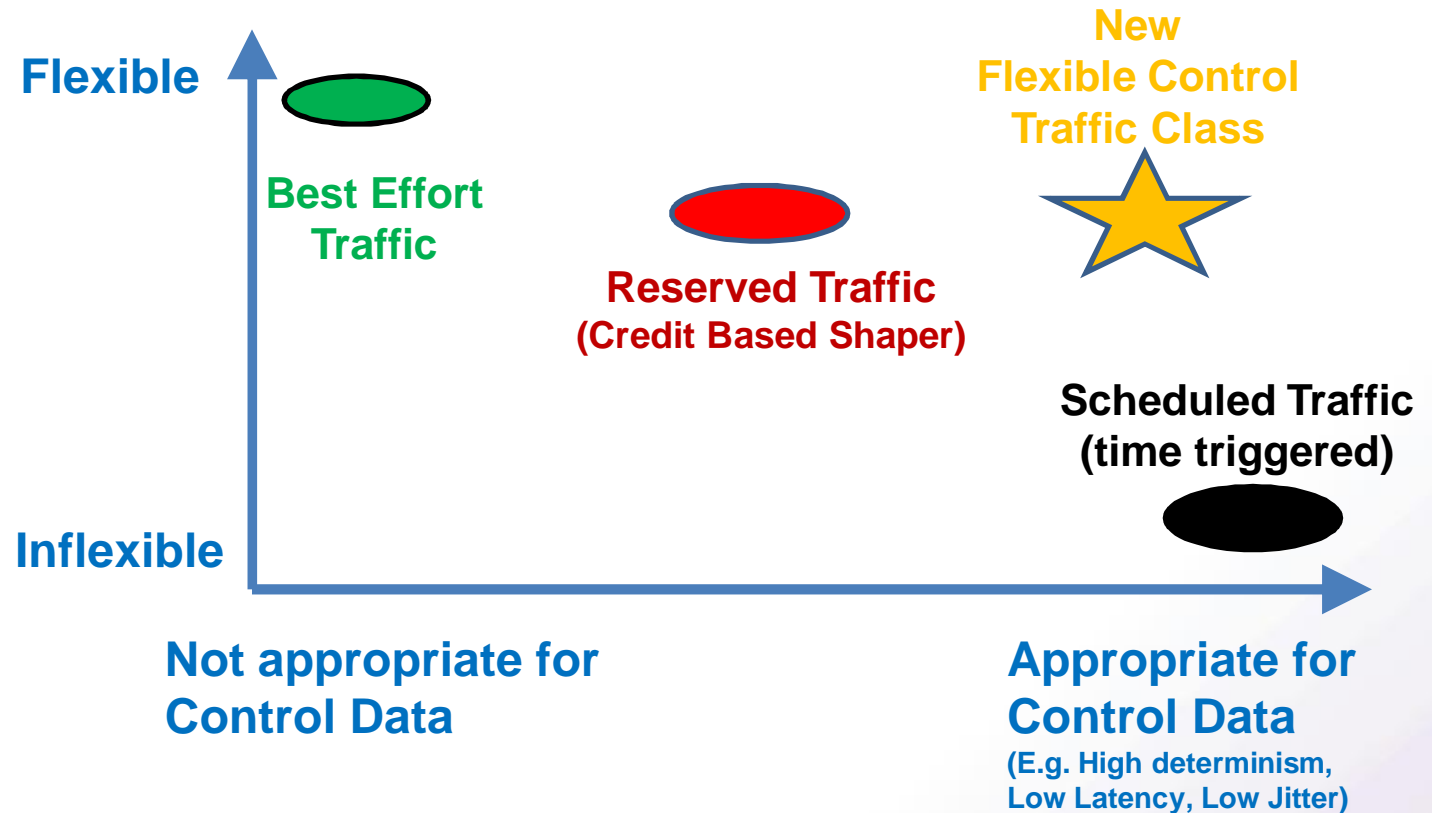
Automotive Requirements for a Flexible Control Traffic Class




Markus Jochim (General Motors)

- 02/05/2014 -

- **Scope & Structure of the Presentation**
- **Part I: Identified Preferences & Requirements**
- **Part II: Discussion Slides**
 - ❖ Category I:
Timing Characteristics of Periodic & Event based Traffic
 - ❖ Category II:
Safety Related Characteristics (Ingress Policing)
 - ❖ Category III:
Establishing Bandwidth & Timing Guarantees (Reservations, SRP)



- IEEE 802.1 TSN is currently working on proposals for additional traffic types with the desired properties: **Flexible AND Appropriate for Control Data** 
- AAA₂C delivers input (= requirements, desired properties) to IEEE.

Structure of the Presentation

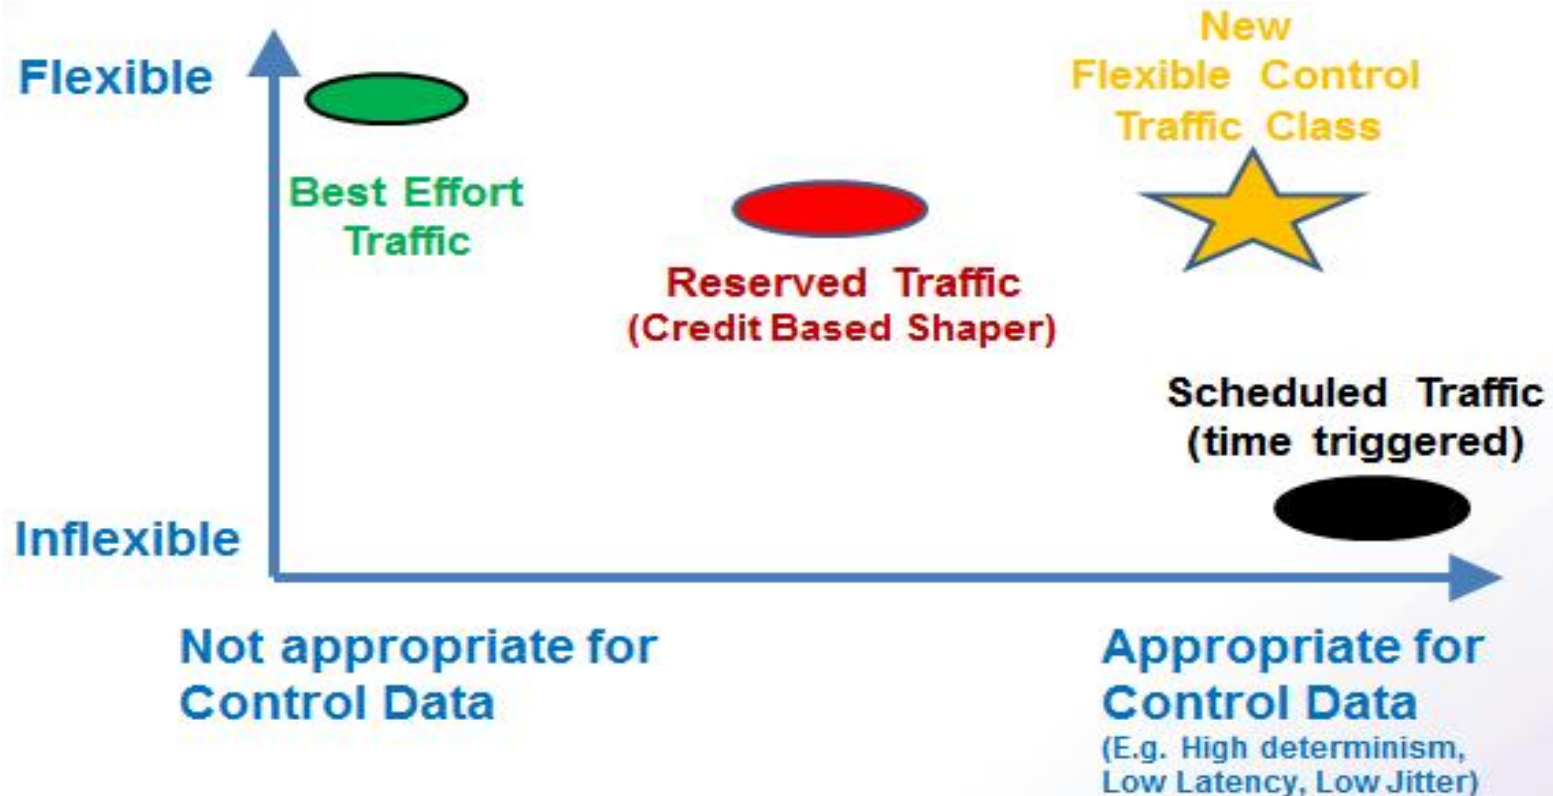
- This slide deck was used in several AAA2C meetings to discuss automotive requirements for a flexible control traffic class.
 - 1) Part I of the presentation summarizes the requirements and preferences that have been identified during the AAA2C meeting.
 - 2) Part II consists of the slides that have been used during the AAA2C meetings to facilitate the technical discussions during which the aforementioned requirements and preferences have been identified.

- **Scope & Structure of the Presentation**
- **Part I: Identified Preferences & Requirements**
- **Part II: Discussion Slides**
 - ❖ Category I:
Timing Characteristics of Periodic & Event based Traffic
 - ❖ Category II:
Safety Related Characteristics (Ingress Policing)
 - ❖ Category III:
Establishing Bandwidth & Timing Guarantees (Reservations, SRP)

Results (1/10)



- This part of the presentation (Part I) summarizes the preferences and requirements for a Flexible Automotive Control Traffic Class that have been identified by the AAA2C group.



Results (2/10)



Date	#	Statements / Requirements / Preferences	Reviewed
05/08	S1	The timing requirements for a flexible automotive control traffic class have been set to values that will enable 90% of the automotive control applications. We assume the most challenging 10% of the applications to be covered by TSN's scheduled traffic (= time triggered traffic).	05/23
05/08, 05/23 Note added.	S2	For periodic traffic, the following range of periods shall be supported: Minimum: 5ms, Maximum: 1000 ms. Example of a typical period: 8 ms. Note: Considering S1, we assume that at least 90% of today's control application will not have periods shorter than 5ms. Looking forward, shorter periods (e.g. 1ms, 2.5ms) are be desirable.	05/23, 06/05
05/08	S3	It shall be possible to freely configure the required periods for periodic messages. A traffic class that supports a fixed small number of predefined periods is not considered adequate.	05/23
05/08	S4	The traffic class needs to support max latency guarantees.	05/23

Results (3/10)



Date	#	Statements / Requirements / Preferences	Reviewed
05/08	S5	The required max latency for each periodic and event based message is known at design time. Different messages have different latency requirements.	05/23
05/08	S6	For periodic and event based traffic, the following minimum latency shall be supported by the traffic class: 1 ms or lower over 7 hops. Example of a typical latency requirement: 3 ms, 7 hops.	05/23
05/23, 06/05 Extended.	S7	We assume that it is difficult to tightly control jitter in a switched network. In presence of time stamping mechanisms and a defined maximum latency, a minimum jitter is not an absolute requirement. A small maximum jitter that is determinable for a given topology is however desirable.	06/05 06/20

Results (4/10)



Date	#	Statements / Requirements / Preferences	Reviewed
05/23	S8	<p>The maximum latency that can be guaranteed is a function of the topology, the number of hops and the link speed.</p> <p>For a new traffic class this function should be as simple as possible so that the network can be analyzed with reasonable effort.</p>	06/05
06/27	S9	<p>Stream Reservation: The availability of a Stream Reservation Protocol (SRP) that supports the Flexible Automotive Control Traffic Class is desirable. The SRP shall accept the required periodicity, the required max. latency and the required bandwidth input parameters the characterize the required Quality of Service. If the QoS requirements can be met by the communication system, the SRP shall lock down the required resources.</p>	08/07

Results (5/10)



Date	#	Statements / Requirements / Preferences	Reviewed
06/27	S10	Stream Reservation: Release Streams The SRP for the Flexible Automotive Control Traffic class shall enable a Talker to release a stream and shall free up the network resources that were previously reserved to accommodate the stream.	08/07
06/27	S11	Identified automotive use cases using SRP for control applications require the ability to: a) Reserve sets of engineered control data streams (By issuing a series of reservation requests). b) Switch between different sets of engineered control data streams. (By issuing a series of release and reservation requests). Note: We currently don't see a strong use case for switching between different sets of non-engineered control data streams.	08/07

Results (6/10): Ingress Policing



- The following slides summarize the results of the discussion of the Ingress Policing topic.
- It is recommended to first review the material on “Ingress Policing” in Part II of this presentation, since the Part II material enables the interpretation of the summary of the discussion.
- The main discussion was focused on answering the 3 questions raised in the following two slides of Part II.



Summary of Observations

Which of the following is acceptable / not acceptable?
Desirable / not desirable? Important / less important?

- Should a faulty talker T be completely silenced?
Meaning: Faulty as well as non-faulty streams from T are silenced.
Or should only faulty streams sent by a faulty talker be silenced?
Do we care?
- Is it acceptable if the Ingress Policing turns a non-faulty streams into a faulty stream as long as the stream was send by a faulty talker?
- Is it acceptable if a non-faulty stream sent by a fault free talker becomes faulty? BLUE
What if this can only happen in presence of a moderate babbler?

Green on next slide

Orange on next slide

Blue on next slide

	Per Stream (= Potentially higher number of filters per port)	Per Class (= Small number of filters per port)
Threshold Enforcing	<ul style="list-style-type: none"> A faulty stream send by a faulty talker is not “silenced”. Other streams from faulty / fault free talkers not affected. 	<ul style="list-style-type: none"> A faulty stream send by a faulty talker is not “silenced”. Non-faulty streams send by faulty talkers can become faulty. A fault free stream send by a fault free talker becomes faulty. (Fault propagation. Fault not contained)
Blocking	<ul style="list-style-type: none"> A faulty stream send by faulty talker is “silenced”. Non-faulty streams send by faulty talker are not necessarily silenced. 	<ul style="list-style-type: none"> If a talker exceeds it's configured bandwidth limit, the faulty talker is “silenced”. In presence of a moderate babbler, a fault free stream send by a fault free talker can become faulty. (Fault propagation. Fault not contained) Faulty streams send by a faulty talker are not necessarily silenced.

Results (7/10): Ingress Policing



- *Question 1: Should a faulty talker T be completely silenced? (*1)*
Meaning: Faulty as well as non-faulty streams from T are silenced.
Or should only faulty streams sent by a faulty talker be silenced?

Summary of the discussion:

- There is value in keeping the mechanism simple by blocking all streams sent by a faulty talker.

Reason: It will typically simplify the design of a fault tolerant system if fail silent behavior of a faulty component can be guaranteed. The fact that a faulty talker is present in the system needs to be addressed at the system level.

- Blocking all streams sent by a fault talker will also solve the issue that *Question 2* raises by preventing situations where Ingress Policing can turn a non-faulty stream sent by a faulty talker into a faulty stream.

(*1): I was pointed out during the discussion that “silencing a talker” or “silencing a stream” are misleading terms, since the discussion focused on blocking either all streams or some streams at the ingress port of the switch the talker is connected to.

Results (8/10): Ingress Policing



- *Question 3: Is it acceptable if a non-faulty stream sent by a fault free talker becomes faulty?*

What if this can only happen in presence of a moderate babbler?

Summary of the discussion:

- Ingress policing should ensure that a faulty stream sent by a faulty talker can not turn a non-faulty stream sent by a fault free talker into a faulty stream. (Independently of whether the aforementioned faulty talker is a moderate babbler or not)
- Based on the answers to questions 1, 2 and 3, the combination “Per Stream x Blocking” (see Part II of the presentation) is the preferred combination.
- The combination “Per Stream x Blocking” should be configureable: When the Filter observes the violation of a threshold on a “per stream” basis, the blocking that follows can be configured to be “per stream blocking” or “per class blocking”.

Conflicting requirements / Limitations:

- The AAA2C group discussed preferences and requirements from a user's perspective.
- The group did not discuss potential implementations and is aware of the fact that some of the requirements may be conflicting requirements.
- In case some of the AAA2C requirements are competing requirements that cannot all be fulfilled by one single traffic class, the next slide gives an indication which limitations of an implementation the AAA2C group would be more / less willing to accept.

Willingness to accept different types of limitations:

Higher
Willingness



- *Max. link utilization not as good as it could be.*
- *Minimum latency is not as good as it could be.
(Related to Req. S6)*
- *Latency turns into a rather complex function of parameters like “Topology, Number of hops, etc.”
Tools are required for analyzing latencies.
(Related to Req. S8)*
- *Latency guarantees may occasionally be violated.
(Related to Req. S4)*

Lower
Willingness

- **Scope & Structure of the Presentation**
- **Part I: Identified Preferences & Requirements**
- **Part II: Discussion Slides**
 - ❖ Category I:
Timing Characteristics of Periodic & Event based Traffic
 - ❖ Category II:
Safety Related Characteristics (Ingress Policing)
 - ❖ Category III:
Establishing Bandwidth & Timing Guarantees (Reservations, SRP)

Discussion Slides

The slides in Part II of the presentation contain questions, statements and observations that have been used during the AAA2C working meetings to stimulate discussions on desirable properties / requirements for a flexible automotive control traffic class. The slides in Part II are not intended to capture discussion results.

- **Scope & Structure of the Presentation**
- **Part I: Identified Preferences & Requirements**
- **Part II: Discussion Slides**
 - ❖ **Category I:**
Timing Characteristics of Periodic & Event based Traffic
 - ❖ **Category II:**
Safety Related Characteristics (Ingress Policing)
 - ❖ **Category III:**
Establishing Bandwidth & Timing Guarantees (Reservations, SRP)

Timing Characteristics of Periodic & Event based Traffic (1/6)

1) Periodic Traffic: “Supported Periods”

- a) What is the range of periods that should be supported by the traffic class? Minimum? Maximum? Typical?
- b) Granularity:
 - i. Is it sufficient if the communication system supports a small number (e.g. 3 or 4) of fixed periods?
 - ii. Or do we need to be able to define a custom set of periods for a given system at design time?
 - iii. Or maybe a specific period for each individual periodic message?

Timing Characteristics of Periodic & Event based Traffic (2/6)

2) Periodic & Event based Traffic: “Latency Guarantees”

Assumptions:

- *We always need max latency guarantees!?*
- *Required latencies for each flow are known at design time!?*

a) What are the max latency requirements (values)?

Background information:

- *Automotive latency requirements as currently discussed within IEEE: 100 μ s over 5 AVB hops. (See slide in backup section)*
- *AVB 1: Class A: 2 ms over 7 hops
Class B: 50 ms over 7 hops* } *Topology dependencies ?*

b) For what “size” of network (number of hops)?

c) Maximum payload size?

Timing Characteristics of Periodic & Event based Traffic (3/6)

3) Periodic Traffic: “Jitter Guarantees”

- a) Do we need a max. jitter guarantee or is a max latency guarantee sufficient?
(Assumption: Global time available)
- b) If we need a max jitter guarantees, what is our max jitter requirement? (Quantitative)

Timing Characteristics of Periodic & Event based Traffic (4/6)

4) Event based Traffic (Jitter Guarantees)

Assumptions:

- *Event based Traffic contains sporadic messages as well as bursty traffic.*
Given our focus on control applications:
Messages that we will find on CAN, LIN and FlexRay's dynamic segment.

a) Required jitter guarantees for event based traffic ?

Timing Characteristics of Periodic & Event based Traffic (5/6)

5) Topology Dependencies (Latency / Jitter Guarantees)

Assumptions:

- *A new traffic class that enables flexible automotive control traffic will be introduced into the TSN standard.
(Currently there are three or more proposals under discussion in 802.1 TSN: Peristaltic Shaper, Burst Limiting Shaper, Urgency based Scheduler)*
 - *Depending on the mechanisms that implement the new traffic class, the dependencies between “Latency and Jitter Guarantees” on one hand, and “Topology” on the other hand, may be more or less complex.*
- a) If we need to re-validate latency guarantees every time we extend an Ethernet network / change a topology / add traffic... is that acceptable from a logistic perspective?
 - b) How important is simplicity?
Is it OK if we need to run tools to understand the overall system behavior w.r.t. latency / jitter guarantees ?

Timing Characteristics of Periodic & Event based Traffic (6/6)

A comment / question received from a participant of the IEEE 802.1 TSN Plenary meeting in Geneva (07/2013) after presenting our requirements:

Assuming that some of the AAA2C requirements that have been presented are competing requirements that cannot be all be fulfilled by one single traffic class. Which of the following ones would you be more / less willing to accept?

- 1) *Latency guarantees may occasionally be violated.* (Related to Req. S4)
- 2) *Latency turns into a rather complex function of parameters like “Topology, Number of hops, etc.”
Tools are required for analyzing latencies.* (Related to Req. S8)
- 3) *Minimum latency is not as good as it could be.* (Related to Req. S6)
- 4) *Link utilization not as good as it could be.*

Note: This slide was used to enable the discussion and the sequence 1), 2), 3), 4) does NOT reflect or summarize the result of the discussion.
For a summary of the result see the slide labeled: “Willingness to accept different types of limitation” in Part 1 of this presentation.

- **Scope & Structure of the Presentation**
- **Part I: Identified Preferences & Requirements**
- **Part II: Discussion Slides**
 - ❖ Category I:
Timing Characteristics of Periodic & Event based Traffic
 - ❖ Category II:
Safety Related Characteristics (Ingress Policing)
 - ❖ Category III:
Establishing Bandwidth & Timing Guarantees (Reservations, SRP)

Ingress Policing

Ingress Policing Discussion during AAA2C Face-2-Face Meeting

- Ingress Policing Topic was already discussed in detail during the AAA₂C Face-to-Face Meeting.
- We discussed the Pros and Cons of some alternatives that will be revisited in the following slides:

	Per Stream	Per Class
Threshold Enforcing	1	2
Blocking	3	4

Note:

F2F Presentation from Yong and Markus is available at <https://council.avnu.org/wg/AAA2C/document/index> in the F2F-Meeting Folder under "Agenda item 4: Ingress Policing".

The aforementioned presentation is structured into 2 parts. Regarding the second part, an updated presentation with a more in-depth analysis has meanwhile become available: <http://ieee802.org/1/files/public/docs2013/tsn-jochim-ingress-policing-1113-v2.pdf>.

Proposal how to proceed today...

- 1) Review some of the F2F slides again
 - The F2F was a couple of months ago, so it might help to refresh our memories.
 - Review will be limited to a subset of the slides presented at the F2F.
- 2) Revisit the observations associated with each of the four alternatives:
{Threshold enforcing, Blocking} X {Per Stream Filter, Per Class Filter}
- 3) Reach a conclusion by discussing what type of observed behavior is acceptable / not acceptable or desirable / not desirable.

Proposal:

- AAA2C should focus on what type of system behavior is acceptable / not acceptable rather than on selecting a mechanism / solution.
- IEEE 802.1TSN should then decide on the Ingress Policing mechanism and how to implement Policing to address the identified requirements.

Babbling Idiot Problem



- A faulty talker or switch (= Babbling Idiot)
 - sends too much traffic or
 - sends at the “wrong time”and takes away bandwidth from other streams.

- Bandwidth and latency guarantees of these “other streams” can no longer be guaranteed.
 - ⇒ **They become faulty !**

- The babbling idiot can affect many streams in a network!
The fault effect propagates through the network and can not easily be contained.

Babbling Idiot Problem

Example:

Babbling Idiot: T1

Faulty red stream sends too much data.

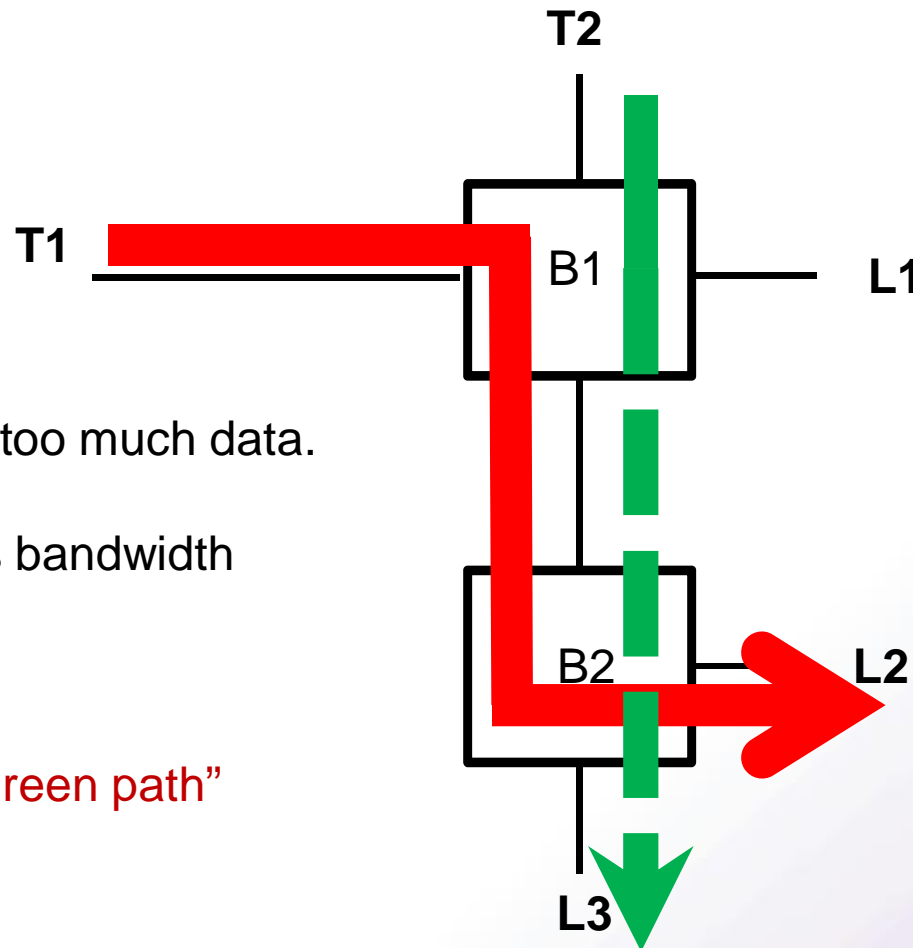
Green stream violates its bandwidth and latency guarantees.

Note:

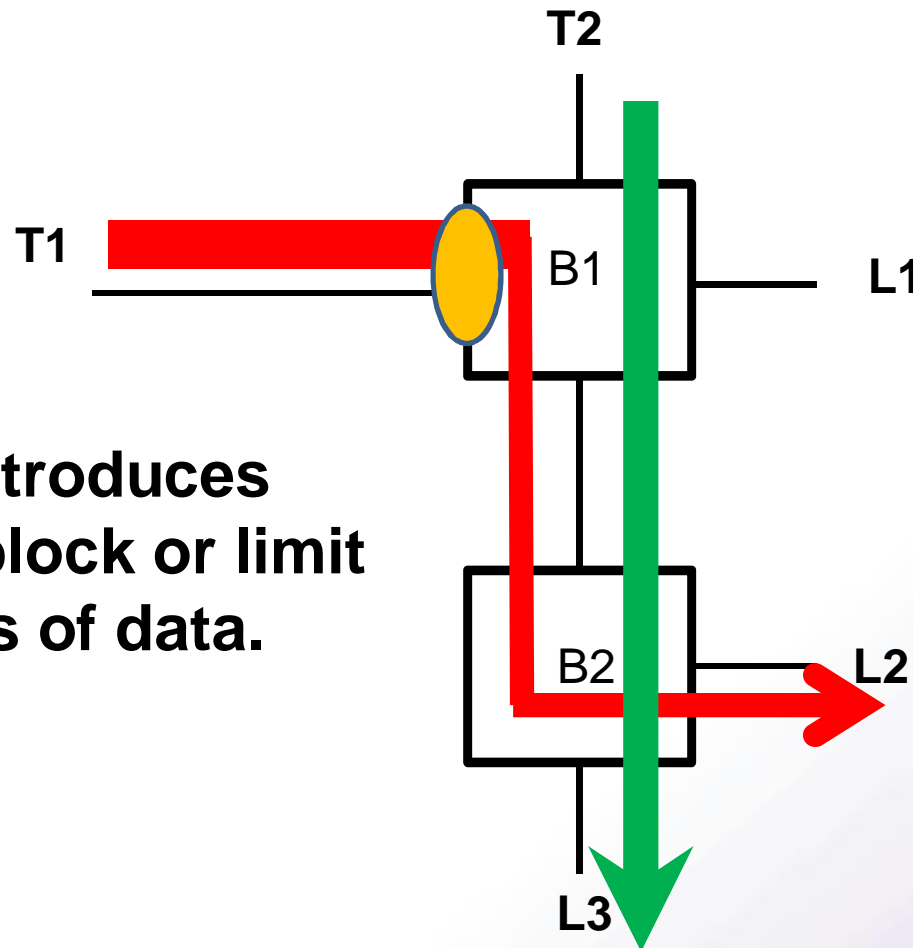
All components on the “green path” are fault free.

But:

Green stream is faulty.



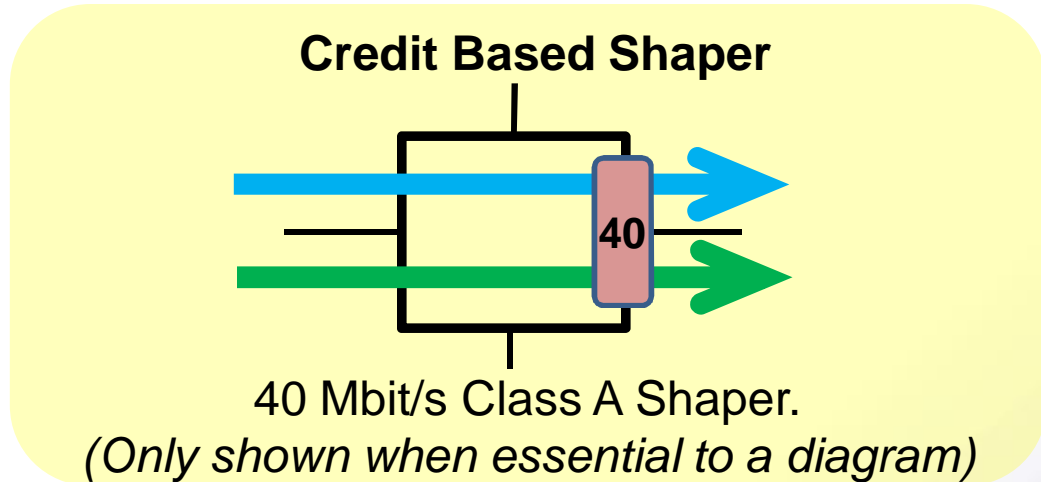
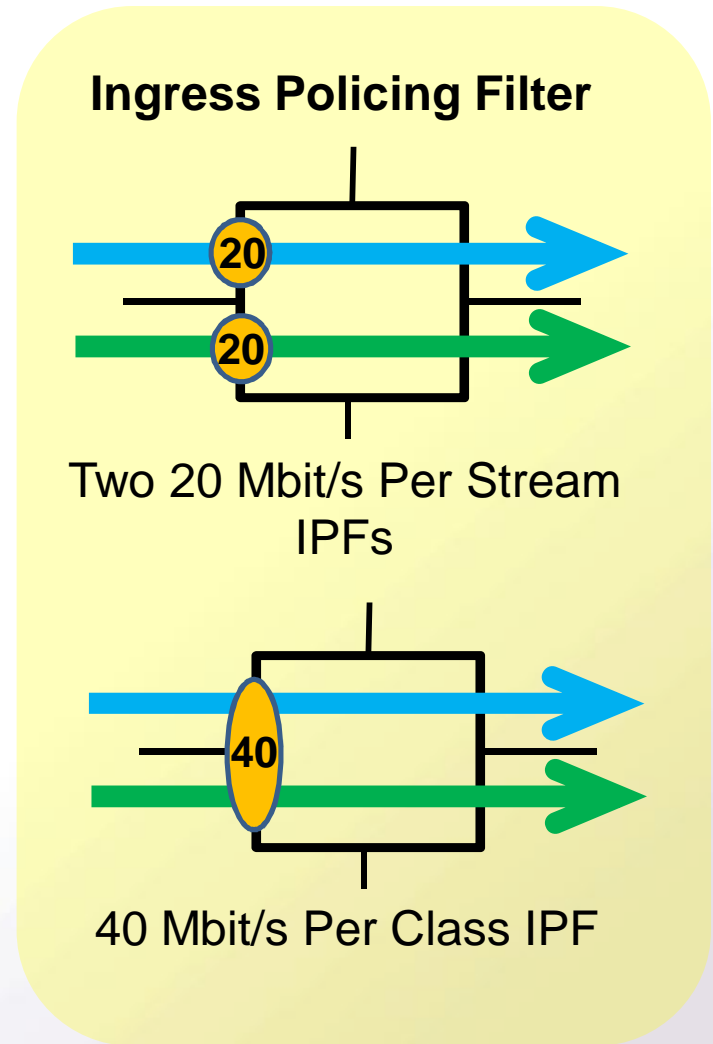

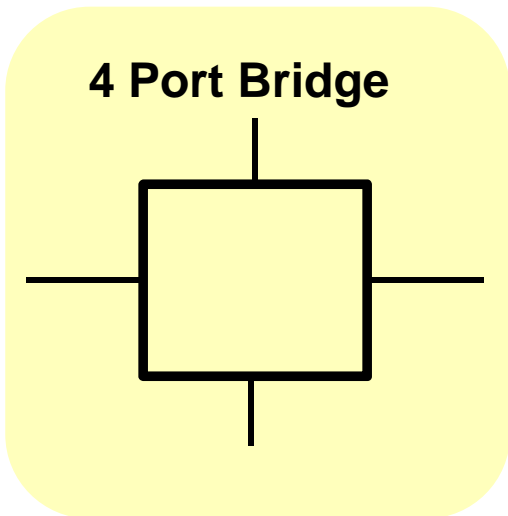
Ingress Policing in a Nutshell



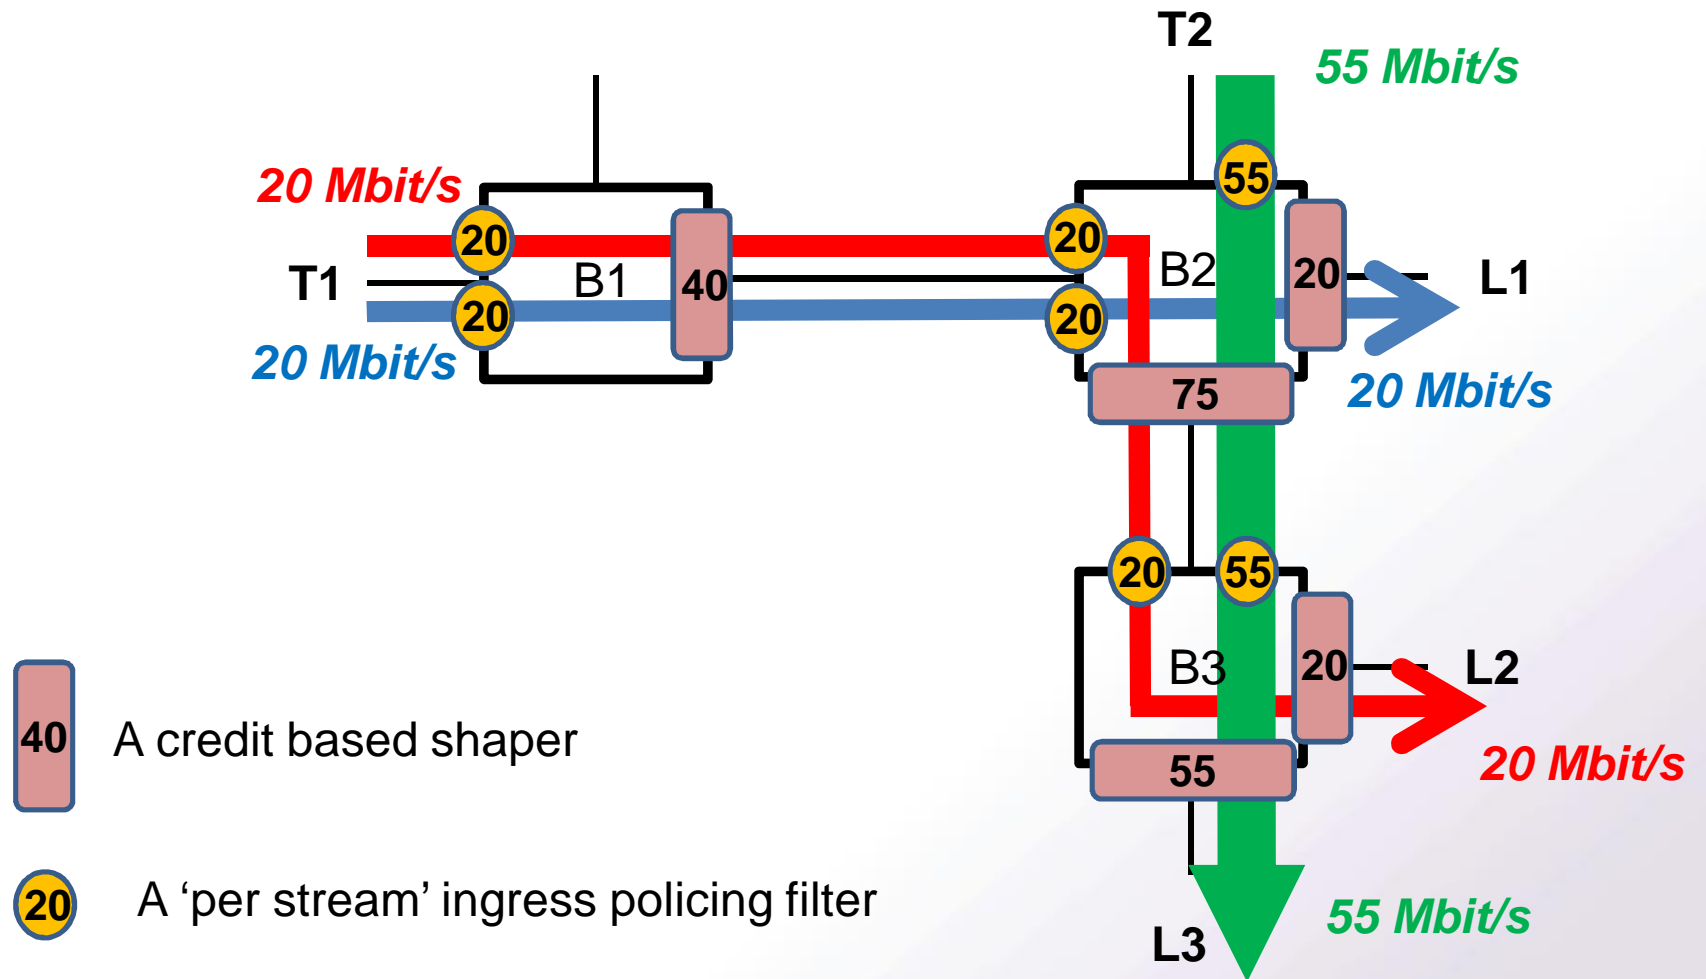
Ingress Policing introduces filters  that will block or limit excessive amounts of data.

Symbols and Abbreviations

- IPF = Ingress Policing Filter
- Talker: T1, T2,... Listener: L1, L2,...



Example: Fault free case...



Example: Per Class X Threshold Enforcing

	Per Stream	Per Class
Threshold Enforcing	1	2
Blocking	3	4


- During the F2F we had a detailed look into all 4 cases
- Let's revisit alternative #2 in detail and then focus on a summary of the results for all four cases.

Per-Class X Threshold Enforcing Filter





Example :

Blue: 20 Mbit/s
Green: 20 Mbit/s



Blue: 20 Mbit/s
Green: 20 Mbit/s

Fault: Blue stream babbles (40 Mbit/s instead of 20 Mbit/s)



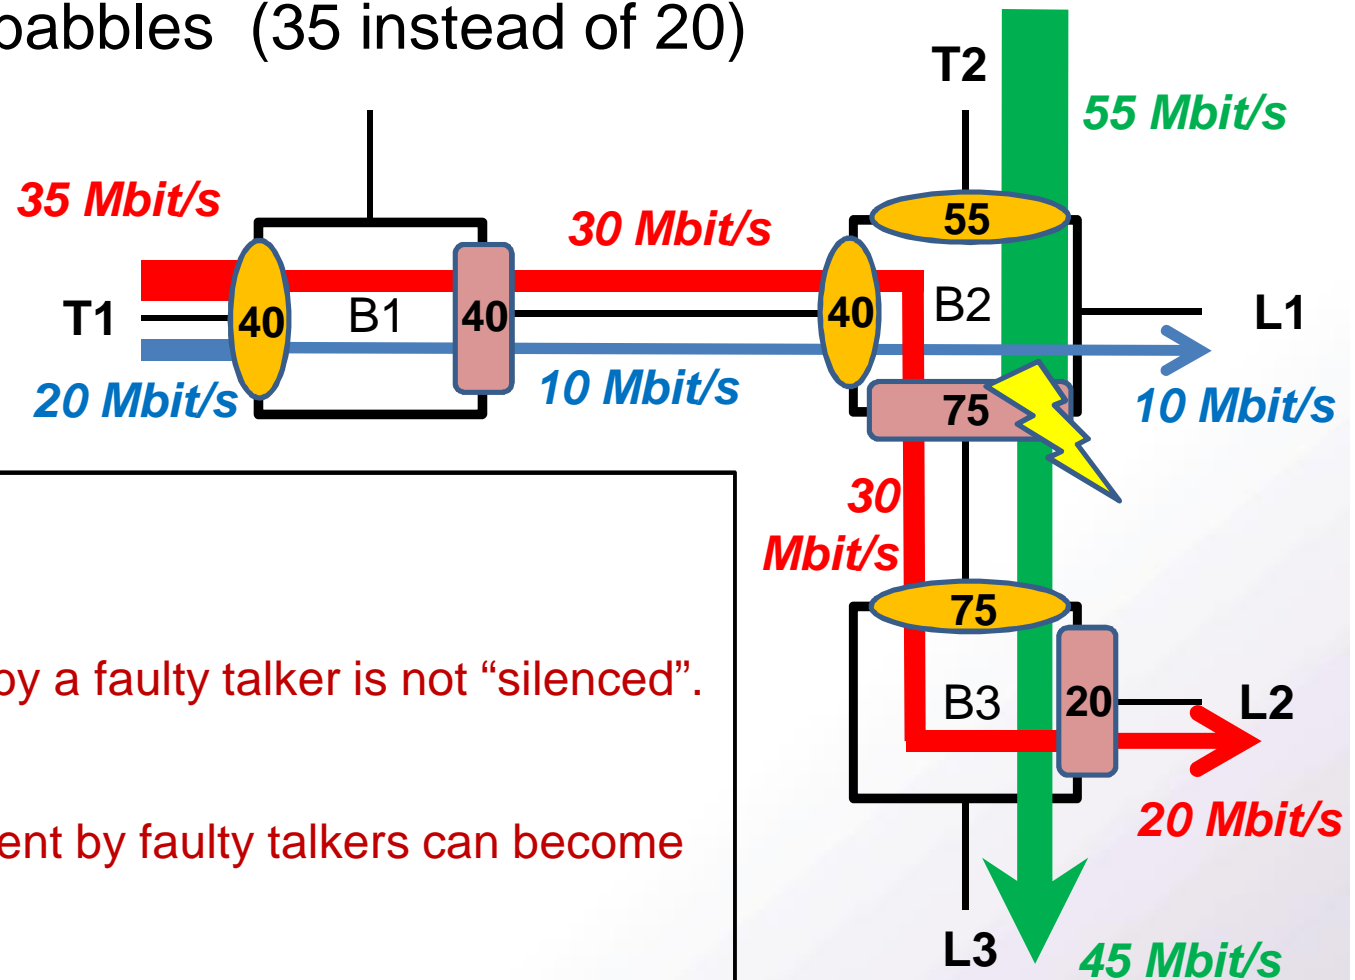
S1: All kinds of behavior (X or Y or anything in between) are possible!

Since a per class ingress policing mechanism is not aware of any streams, it can only discard arbitrary class A frames once the established bandwidth threshold is exceeded. The discarded frames could be blue frames only, or green frames only, or any mix of blue and green frames we can think of.


Per Class X Threshold Enforcing



- Fault: **T1-red** babbles (35 instead of 20)



Observations:

- T1-red:
A faulty stream sent by a faulty talker is not “silenced”.
- T1-blue:
Non-faulty streams sent by faulty talkers can become faulty.
- T2-green:
A fault free stream sent by a fault free talker becomes faulty. (Fault propagation. Fault not contained) 


Note: This diagram shows one out of many different ways of how things could play out. (See statement S1 on previous slide)

Four Alternatives (1/2)


- Before moving on to the summary of the observations from all four alternatives from the F2F meeting, let's remind ourselves what the four alternatives really are.

	Per Stream	Per Class
Threshold Enforcing	1	2
Blocking	3	4

- We already saw the difference between “Per Stream” and “Per Class” IPFs:



Two 20 Mbit/s Per Stream IPFs

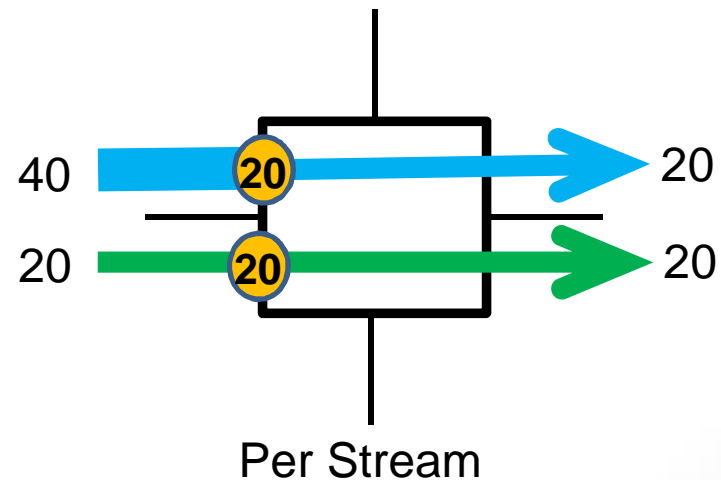



40 Mbit/s Per Class IPF

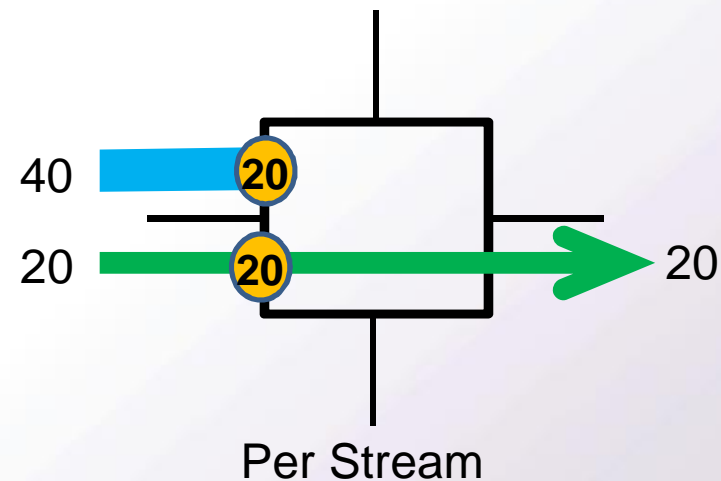
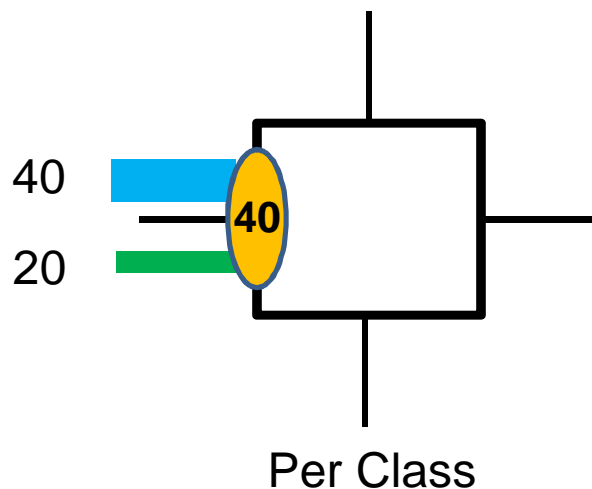
Four Alternatives (2/2)



➤ And we already looked into Threshold Enforcing Filters:



➤ And Blocking Filters are really simple:



Next three slides:


Summary of observations.

With some updates based on a refined analysis.

Per Stream

(= Potentially higher number of filters per port)

Threshold Enforcing



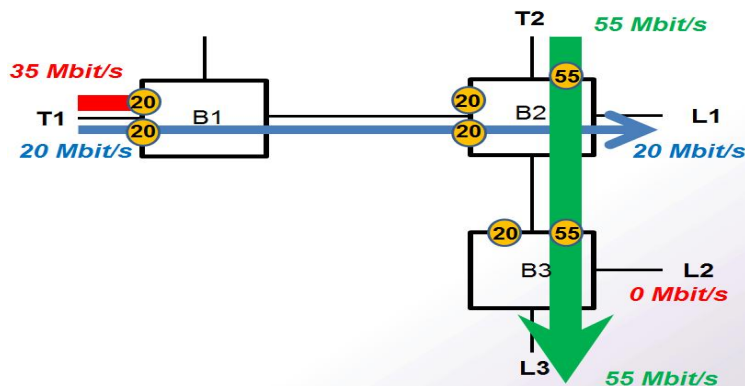
- A faulty stream sent by a faulty talker is not “silenced”.
- Other streams from faulty / fault free talkers not affected.

Per Class


(= Small number of filters per port)

- A faulty stream sent by a faulty talker is not “silenced”.
- Non-faulty streams sent by faulty talkers can become faulty.
- A fault free stream sent by a fault free talker becomes faulty. (Fault propagation. Fault not contained)

Blocking



- A faulty stream sent by a faulty talker is “silenced”.
- Non-faulty streams sent by faulty talker are not necessarily silenced.



- If a talker exceeds its configured bandwidth limit, the faulty talker is “silenced”.
- In presence of a moderate babbler, a fault free stream sent by a fault free talker can become faulty. (Fault propagation. Fault not contained).
- Faulty streams sent by a faulty talker are not necessarily silenced.

Summary of Observations

Which of the following is acceptable / not acceptable?
Desirable / not desirable? Important / less important?

- 1) Should a faulty talker T be completely silenced?
Meaning: Faulty as well as non-faulty streams from T are silenced.

Or should only faulty streams sent by a faulty talker be silenced?

Do we care?

Green
on next
slide

- 2) Is it acceptable if the Ingress Policing turns a non-faulty streams into a faulty stream as long as the stream was send by a faulty talker?

Orange
on next
slide

- 3) Is it acceptable if a non-faulty stream sent by a fault free talker becomes faulty? BLUE

What if this can only happen in presence of a moderate babbler? (*1)


Blue
on next
slide

(*1): For a definition of the term 'moderate babbler' and an analysis of potential fault effects in presence of moderate babblers see:
<http://iee802.org/1/files/public/docs2013/tsn-jochim-ingress-policing-1113-v2.pdf>.

Per Stream

(= Potentially higher number of filters per port)

Threshold Enforcing




- A faulty stream send by a faulty talker is not “silenced”.
- Other streams from faulty / fault free talkers not affected.

Per Class

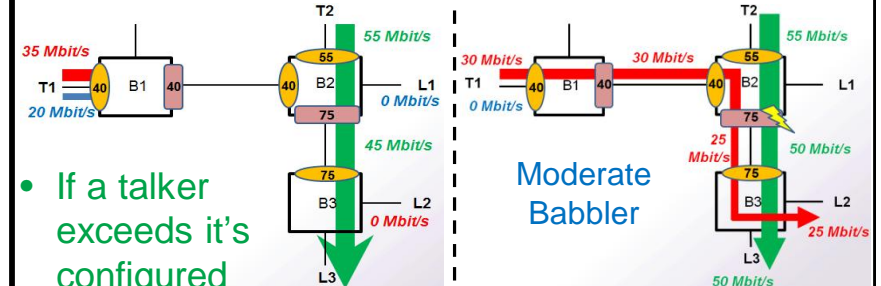
(= Small number of filters per port)

Blocking



- A faulty stream send by faulty talker is “silenced”.
- Non-faulty streams send by faulty talker are not necessarily silenced.

- A faulty stream send by a faulty talker is not “silenced”.
- Non-faulty streams send by faulty talkers can become faulty.
- A fault free stream send by a fault free talker becomes faulty. (Fault propagation. Fault not contained)



- If a talker exceeds it's configured bandwidth limit, the faulty talker is “silenced”.
- In presence of a moderate babbler, a fault free stream send by a fault free talker can become faulty. (Fault propagation. Fault not contained).
- Faulty streams send by a faulty talker are not necessarily silenced.

How fast do Filters need to respond?

- Filter latency:
 - Traffic needs to be observed over some time to decide whether or not a bandwidth threshold is exceeded (= babbling detected)
 - Once babbling is detected, it takes time to switch on the desired response (either threshold enforcing or blocking).
 - May take more / less time, depending on whether or not the switch can trigger a preconfigured action, or interaction with a microcontroller is required.

- Requirements related to filter latency?

- **Scope & Structure of the Presentation**
- **Part I: Identified Preferences & Requirements**
- **Part II: Discussion Slides**
 - ❖ Category I:
Timing Characteristics of Periodic & Event based Traffic
 - ❖ Category II:
Safety Related Characteristics (Ingress Policing)
 - ❖ Category III:
Establishing Bandwidth & Timing Guarantees (Reservations, SRP)

Reservations: Establishing Bandwidth and Timing Guarantees (1/6)

Assumption:

- The assumption is that flexible control traffic can coexist with other traffic classes (BE, reserved traffic, scheduled traffic) on the same network and that some of these traffic classes (reserved traffic, scheduled traffic) will use existing mechanisms (like SRP) to make reservations.
- However, the purpose of the following slides to discuss, which properties static or dynamic reservation mechanisms for the flexible control traffic class should have.

Reservations: Establishing Bandwidth and Timing Guarantees (2/6)

1) **Dynamic Reservations Mechanisms for the Flexible Control Traffic Class**

- Do we need dynamic reservation mechanisms that evaluate control stream reservation requests at runtime (SRP like)?
- What are the parameters for a reservation request: Only bandwidth and required max. latency?
- Is it acceptable if a runtime mechanism denies a reservation request for a new control stream?
- Is that maybe acceptable if we classify streams into “critical / mission critical” and “non-critical” streams?
- Are there emergency situations, where a dynamic mechanism should reconfigure the system to drop an existing stream (control stream or other stream) in favor of a new critical control stream?
- Does our picture change, when we consider that other traffic classes may have their own reservation mechanisms?

Reservations: Establishing Bandwidth and Timing Guarantees (3/6)

2) **Static Reservations Mechanisms for the Flexible Control Traffic Class**

- Do we prefer a static reservation mechanisms. Would an offline tool evaluate control stream reservation requests at design time?
- Is that problematic from a logistic perspective?
(Many applications. Required central database. Complex change management processes. Complicated by the fact that multiple automotive suppliers designing subsystems. Etc.)

(Continued on next slide)

Reservations: Establishing Bandwidth and Timing Guarantees (4/6)

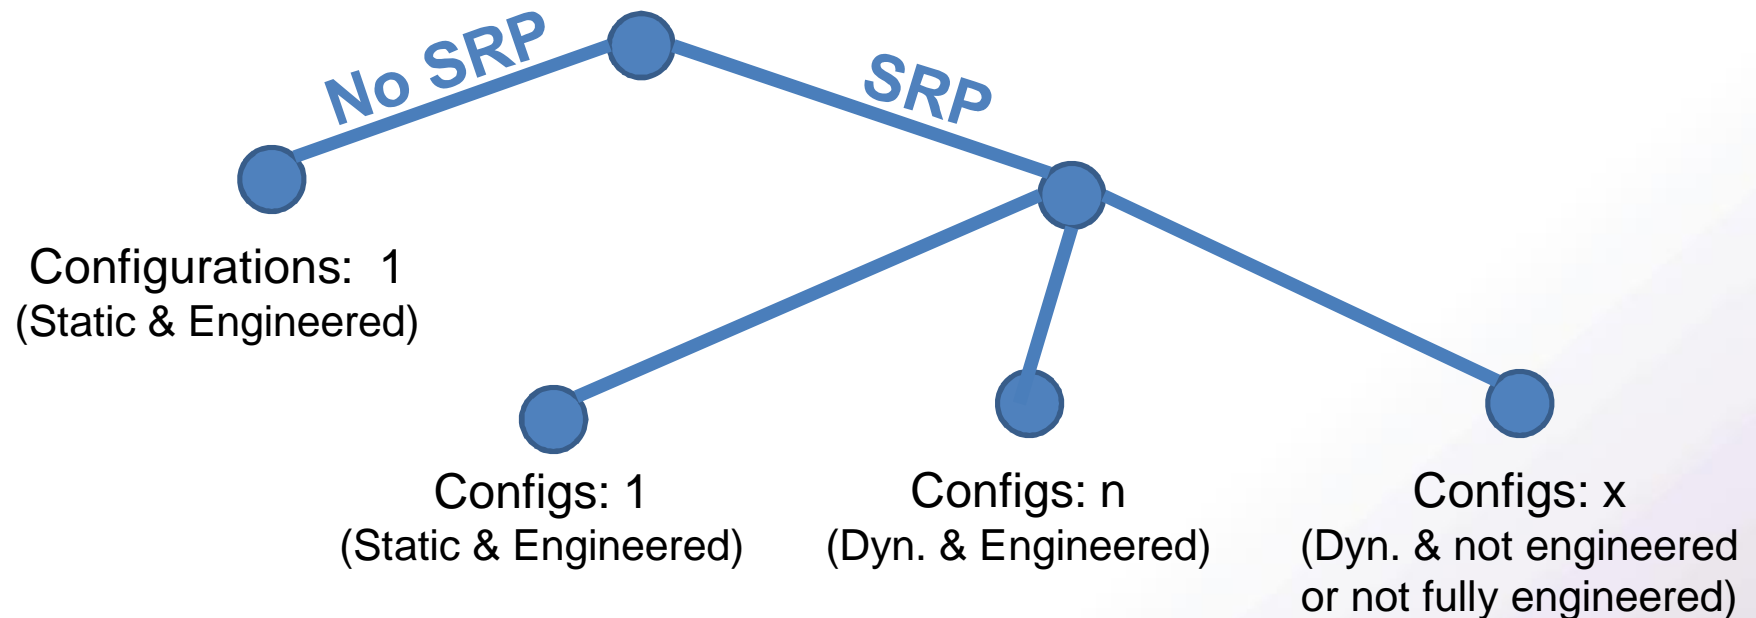
(Continued from previous slide)

- Are there situations (e.g. emergency situations), where we would reconfigure the system at runtime, to drop an existing stream (control stream or other stream) in favor of a new critical control stream?

Assumption: We planned the mode change (e.g. from “normal operation” to “emergency”) at design time and planned which stream will be dropped to free up bandwidth, reduce latency etc.

- Does our picture change, when we consider that other traffic classes may have their own reservation mechanisms?

Reservations: Establishing Bandwidth and Timing Guarantees (5/6)



n: known at design time

x: not known at design time

Reservations: Establishing Bandwidth and Timing Guarantees (6/6)

	No SRP	SRP	SRP	SRP
	<i>Configs: 1</i>	<i>Configs: 1</i>	<i>Configs: n</i>	<i>Configs: x</i>
Characteristics	<ul style="list-style-type: none"> ➤ Static & Engineered ➤ No reconf. during operation. 	<ul style="list-style-type: none"> ➤ Static & Engineered ➤ No reconf. during operation. 	<ul style="list-style-type: none"> ➤ Dyn. & Engineered ➤ n "prepared" configs. ➤ n typically small (e.g. 2 or 3) 	<ul style="list-style-type: none"> ➤ Dyn. & not or not fully engineered ➤ x configs. ➤ x unknown at design time
Pro & Con	<ul style="list-style-type: none"> + Simple (V&V) + Quick startup + No SRP overhead - Adding / modifying streams at design time requires changes in multiple devices. 	<ul style="list-style-type: none"> + Simple (V&V) + Vehicle Developm.: Adding / modifying streams at design time is simple. - Startup time ??? - SRP overhead 	<ul style="list-style-type: none"> + Enables use cases like "switch to vehicle programming mode" or "switch to emergency / limp home mode" - More complex (V&V) due to mode changes - Startup time ??? - SRP overhead. 	<ul style="list-style-type: none"> + Enables use cases like "switch to vehicle programming mode" or "switch to emergency / limp home mode" + Enables dynamic decisions / algos. - Very complex (V&V) - Startup time ??? - SRP overhead.

Backup

The content below was copied from
“QoS requirements for Automotive backbone systems”
Yong Kim (Broadcom), Junichi Takeuchi (Renesas), Masa Nakamura (Envital)
<http://www.ieee802.org/1/files/public/docs2011/new-avb-nakamura-automotive-backbone-requirements-0907-v02.pdf>

QoS requirements for automotive control data class

Performance requirements for automotive control data class

- Maximum latency: 100 us / 5 AVB hops
 - Guaranteed latency
 - Topology independent
 - Automotive control data class to have higher priority than SR classes
 - Maximum 2 priority classes (e.g. Control data class and SR class A)

Preconditions for performance requirements

- Network type: Dedicated network in a vehicle
- Network attributes
 - Maximum AVB hop count: 7
 - Maximum number of nodes (bridged end station & end stations): 32
 - Maximum cable length: 24 m
 - Maximum end-to-end cable length: 30 m
- Automotive control data class attributes
 - Maximum data size (payload size): 128 bytes @FE ~ 256 bytes @GE
 - Maximum number of simultaneous transmission: 8 @FE ~ 32 @GE
 - Transmission period: 500 us
- Payload size for other/lower traffic classes: 256 bytes @FE ~ 1500 bytes @GE

These are our best estimates derived from multiple assumptions of the current and future automotive applications.