Missing Station Capabilities for Resource Engineering

Johannes Specht (Self)

Contributors:

Norman Finn (Huawei)

Günter Steindl (Siemens AG)

Stephan Kehrer (Hirschmann)

Rodney Cummings (National Instruments)

Overview

Introduction

This slide set is for discussion towards standardizing device capabilities, including resource limitations, which is currently not or insufficiently addressed in IEEE 802.1 Standards.

Contents

1. Problem Description
2. Possible Solutions in IEEE 802.1 Standards
3. On Unified Models in IEEE 802.1 Standards

Problem Description

Problem Statement

The Problem

Capabilities and resource limitations of Bridges and End Stations need to be available in a manner suitable for online and offline network engineering by machines (and humans).

Example Capabilities and Resource Limitations

Capability/Limitation	IEEE 802.1 Stds References
Guaranteed maximum gate control list size	12.29.1.5 of IEEE Std 802.1Q-2018
Supported transmission selection algorithms	None
Guaranteed FDB capacity	None
Guaranteed frame buffer capacities	None
Guaranteed end station transmission performance (e.g., imagine many minimum sized frames need to be sent back-to-back)	None
…	…

Why, and where?

There are application areas that require guarantees before network installations or modifications become effective

- Testing/measuring if it works (or not) after installation or modification is inacceptable from an economic and/or technical point of view, …
- … or, in simple words, too late!

Cases (and there are more)

1. When developing fully engineered deterministic networks from scratch
2. When reconfiguring existing fully engineered deterministic networks
   1. When developing extensions for existing deterministic networks
   2. When integrating extensions into existing deterministic networks

Next slides!

Reconfiguring existing deterministic networks

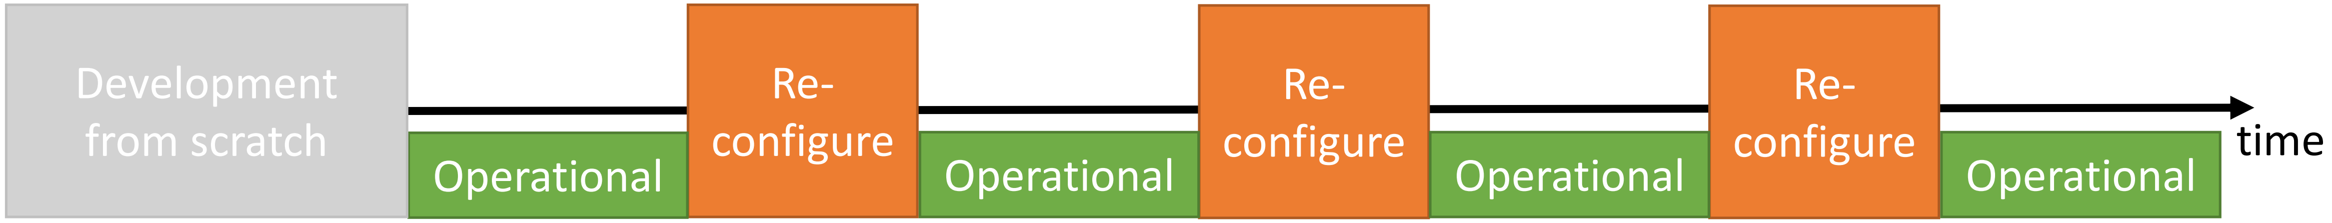

Re-configuration

- Adding, removing, or changing physical network segments/plugging things together differently
- Adding, removing, or changing streams, stream parameters, and paths
- …

Environmental Goals and Constraints

- Avoid long down-times/keep the network operational (cost, availability, etc.)
- Ensure that a re-configuration <u>will</u> work before performing it
- Not necessarily a network engineer available …
- .. And even if:
  - Re-configuration algorithms (e.g., CNC) are necessary due to the complexity of deterministic networks
  - Re-configuration algorithms require machine readable presentations of <u>device capabilities and resource limitations</u> to ensure success

Developing Extensions for Existing Deterministic Networks

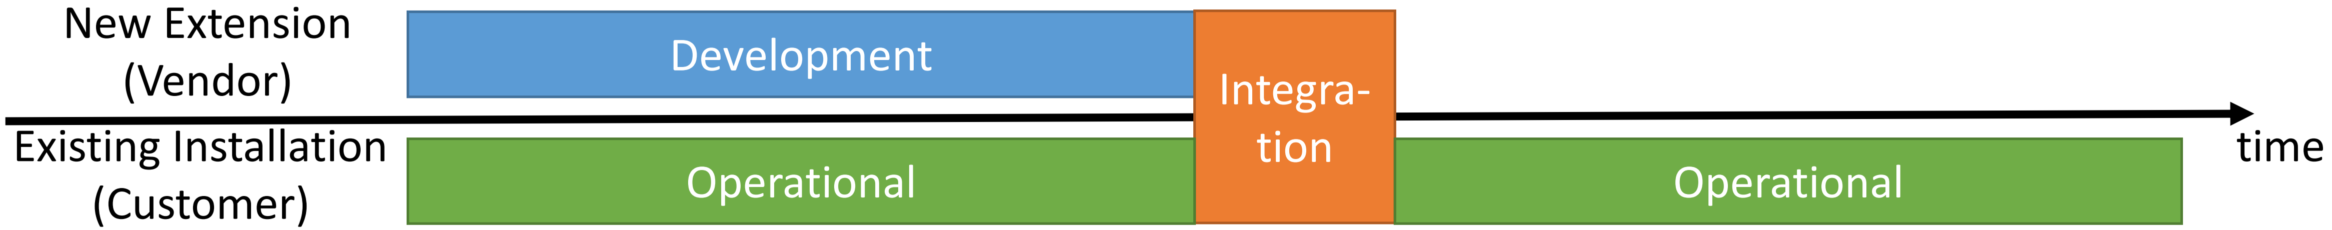

Example from Industrial Automation: Extension of Production Capabilities

Integration should be efficient, without long shutdowns/debugging and alike

- Requires known device capabilities and resource limitations during development/before integration.
  - **Allows preventing/working around bottlenecks during development**
    Example: Planning a non-redundant stream path clockwise instead of counter-clockwise in an existing ring topology to avoid FDB overflows.
  - **Allows cost and time efficient upgrades in existing installations**
    Example: Replacing a single Bridge known to provide insufficient frame buffers.
- **On-site trial & error iterations during integration and manual inspection should be avoided!**

Integrating Extensions into Existing Deterministic Networks

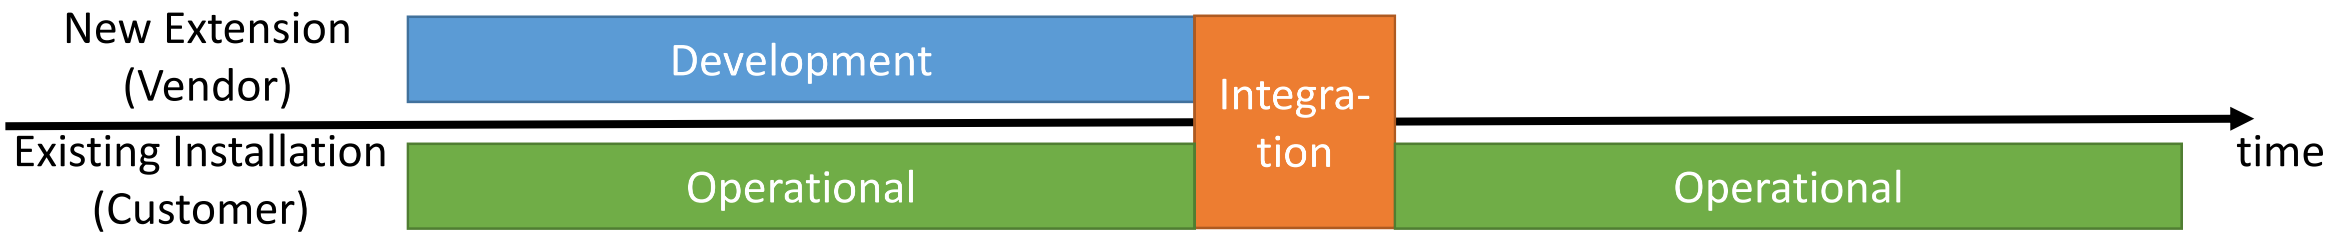

Extending networks implies extending or re-computing network configurations

- Several computational demanding problems (e.g., optimal redundant paths, accurate delay bounds, TDM schedules).
- Some options:
  1. Let a CNC compute without time-consuming random/non-converging trial & error iterations. **Avoid the following**:
     a) CNC computes network configuration 1 (takes some time), Bridges **A** and **C** reject it for reasons unknown
     b) CNC computes network configuration 2 (takes some time), now Bridges **A** and **D** reject it for reasons unknown
     c) CNC computes network configuration 3 (takes some time), now Bridges **B** and **C** reject it for reasons unknown
     d) …

     → **Information on device capabilities and resource limitations needed before CNC computation.**
  2. Calculate new network configurations during development, provide these to simple CNCs at integration time

     → **Information on device capabilities and resource limitations needed before development.**
  3. Guidance from information collected during development for heuristics in clever CNCs

     → **Information on device capabilities and resource limitations needed before development and CNC computation.**

Possible Solutions in IEEE 802.1 Standards

Conformance Statements

Capability/Limitation	IEEE 802.1 Stds References
Guaranteed maximum gate control list size	12.29.1.5 of IEEE Std 802.1Q-2018
Supported transmission selection algorithms	None
Guaranteed FDB capacity	None
Guaranteed frame buffer capacities	None
Guaranteed end station transmission performance (e.g., minimum frames sent back-to-back)	None
…	…

Conformance Statements in technical core Standards (e.g., 802.1Q)

Could be an option (e.g., “A TSN End Station shall be capable to transmit frames of any size back-to-back”), but a limited one:

- Risk to maximize/oversize device requirements all across the board, although …
- … requirements differ significantly between application areas/markets (i.e., no “one size fits all”)

Profiles

Capability/Limitation	IEEE 802.1 Stds References
Guaranteed maximum gate control list size	12.29.1.5 of IEEE Std 802.1Q-2018
Supported transmission selection algorithms	None
Guaranteed FDB capacity	None
Guaranteed frame buffer capacities	None
Guaranteed end station transmission performance (e.g., minimum frames sent back-to-back)	None
…	…

802.1 Profile Projects

Similar to the previous slide, with the following additions:

- Resolution is higher (e.g., different markets), which may mitigate the risk for oversizing to certain extent.
- Adding requirements for elements not standardized in a core Standard (e.g., buffer capacities) is an issue, because defining such elements is typically beyond the scope of profile projects and needs to be done in technical core Standards first.

Unified Models

Capability/Limitation	IEEE 802.1 Stds References
Guaranteed maximum gate control list size	12.29.1.5 of IEEE Std 802.1Q-2018
Supported transmission selection algorithms	None
Guaranteed FDB capacity	None
Guaranteed frame buffer capacities	None
Guaranteed end station transmission performance (e.g., minimum frames sent back-to-back)	None
…	…

Unified Models for Device Capabilities and Resource Limits in core 802.1 Standards (e.g., IEEE Std 802.1Q)

More versatile, if available via management/YANG/MIB (CNC) and for offline engineering. However:

- Providing device capabilities and resource limitations may not always be desired.
  But providing these is already a requirement for certain application areas/markets, then …
- … finding **the right abstraction level** appears important:
  - Overpessimism/cuts in device capabilities and resource limitations can be assumed to be low in generic and detailed models, but higher complexity can be expected.
  - The opposite appears to be true for rather simple models.
  - One option is to go for generic and detailed models that allow reduction to simpler models when used.

On Unified Models in 802.1 Standards

Possible managed objects for an initial FDB capacity model.

Three(!) first models for buffer capacities at different abstraction levels.

The models are intended to give an impression of what is meant by “Unified Models”, and to encourage discussion.

A first model for FDB Capacities

Some possible managed objects:

1. Number of FDB entries currently occupied.
2. Best-case number of available FDB entries (e.g. the FDB is a CAM, or the addresses perfectly fit the hash buckets).
3. Worst-case number of available FDB entries before an attempt to store a new entry can fail due to lack of resources (e.g. hash bucket full).
4. Worst-case number of available FDB entries before forwarding may be delayed for some destination addresses due to increased lookup time.
5. Number of failed attempts to add a new FDB entry.
6. Boolean: Some frames may encounter forwarding delays because of the number or distribution of entries in the FDB.

A (first) first model for Buffer Capacities

Managed object specifying the available buffers as a function of time, link speed and traffic class (independent from the frame size).

Illustration

- $T_{A,1}$ ms @ R_1 Mbit/s for class *A*, $T_{B,1}$ ms @ R_1 Mbit/s for class *B*, …
- $T_{A,2}$ ms @ R_2 Mbit/s for class *A*, $T_{B,2}$ ms @ R_2 Mbit/s for class *B*, …

Time is to be interpreted as uninterrupted transmission by the associate class at the given link speed.

A (second) first model for Buffer Capacities

Each system has some number of buffer pools, each with a maximum number of frames and a maximum number of wire-time bits. (Thus, this only works for Ethernet.) There are four variables per pool for currently-used and available values, times frames and bits.

Each class-of-service queue on each port is associated with one buffer pool.

A (third) first model for Buffer Capacities

The model of each system may have managed objects for the following:

- A frame buffer pool, where each frame buffer is characterized by capacity [Bytes], chunk size[Bytes], and overhead per chunk [Bytes]. For example, a frame buffer may have 128B chunks or 2KB chunks.
- A queue buffer pool, where each queue buffer is characterized by an integer maximum number of frames per queue.
- Separation between frame buffers and queue buffers (that hold pointers to frame buffers) allows both being independent from each other in terms of resolutions and scopes. For example, a central frame buffer with per port per class queue buffers.
- A linkage structure for association of traffic classes per transmission ports with the former for the following:
  - Define which associations are supported by a system (i.e., read-only).
  - Define default associations, and permit constraint user-configured associations.

Extensions

- Provide parameters for fixed allocation (e.g., reserving a certain share of a frame buffer for one per-port class exclusively) - extending the model to support this appears simple and often required (e.g., protect against best effort stealing all buffers).
- Reception-port → buffer and buffer → buffer associations (e.g., for line-card Bridges), but maybe this can hidden.

How to Proceed?

Current Situation

- Industrial Automation needs unified models, rather sooner than later (60802 timing), and
- other markets (e.g., Automotive) would likely benefit from such models.

If IEEE WG 802.1 can reach consensus that unified models are needed (?), what would be the right vehicle?

- Unified models can't be provided in profiles (e.g., 60802), as explained earlier. Technical 802.1 core standards are the appropriate location.
- Could this be done in running projects, e.g. IEEE P802.1Qdj?

Scope of the project:
This amendment specifies procedures, interfaces, and managed objects to enhance the three models of 'Time-Sensitive Networking (TSN) configuration'. It specifies enhancements to the User/Network Interface (UNI) to include new capabilities to support bridges and end stations in order to extend the configuration capability. This amendment preserves the existing separation between configuration models and protocol specifications. This amendment also addresses errors and omissions in the description of existing functionality.

Thank you for your Attention!

Questions, Opinions, Ideas?

Johannes Specht
Dipl.-Inform. (FH)
GERMANY
johannes.specht.standards@gmail.com

Backup

Backup: Datasheets

Capability/Limitation	IEEE 802.1 Stds References
Guaranteed maximum gate control list size	12.29.1.5 of IEEE Std 802.1Q-2018
Supported transmission selection algorithms	None
Guaranteed frame buffer capacities	None
Guaranteed FDB capacity	None
Guaranteed end station transmission performance (e.g., minimum frames sent back-to-back)	None
…	…

Device-Vendor Specific PDF Datasheets

- Unreadable by machines.
  Example: CNCs.
- Even if these would be translated into a machine-readable form, developing configuration algorithms can become hard to impossible from a cost and technical point of view.
  Example: Different devices from different vendors in the same network, all with different machine-readable models.
- May not contain all relevant information.
  Example: It is sometimes more relevant to know the worst-case guaranteed capacity of an FDB without performance penalties instead of the best-case capacity.
- May not be available at all.

Backup: FRER and Retransmission

Capability/Limitation	IEEE 802.1 Stds References
Guaranteed maximum gate control list size	12.29.1.5 of IEEE Std 802.1Q-2018
Supported transmission selection algorithms	None
Guaranteed frame buffer capacities	None
Guaranteed FDB capacity	None
Guaranteed end station transmission performance (e.g., minimum frames sent back-to-back)	None
…	…

Mitigating Occasional Buffer Overruns by Redundancy (e.g., FRER)

Redundancy …

→more frames …

→more buffer utilization …

→more buffer overruns.

Mitigating Occasional Buffer Overruns by Retransmission (e.g., TCP)

- Typically won't satisfy the end-to-end delay requirements at all (i.e., no time for extra roundtrips).
- Determining guaranteed end-to-end delay bounds is a problem (too much dynamics).