

AVB for low latency / industrial networks:

Redundancy for fault-tolerance and AVB – Overview of the simultaneous multi-path proposal



Oliver Kleineberg, Hirschmann Automation & Control

(oliver.kleineberg@belden.com)

IEEE 802 Plenary Meeting, March 2012, Kailua Kona, USA

Agenda:

- **Assumptions and Requirements**
- **Overview of the Proposal**
- **Example application scenarios and benefits**
- **Conclusion**

Assumptions and Requirements

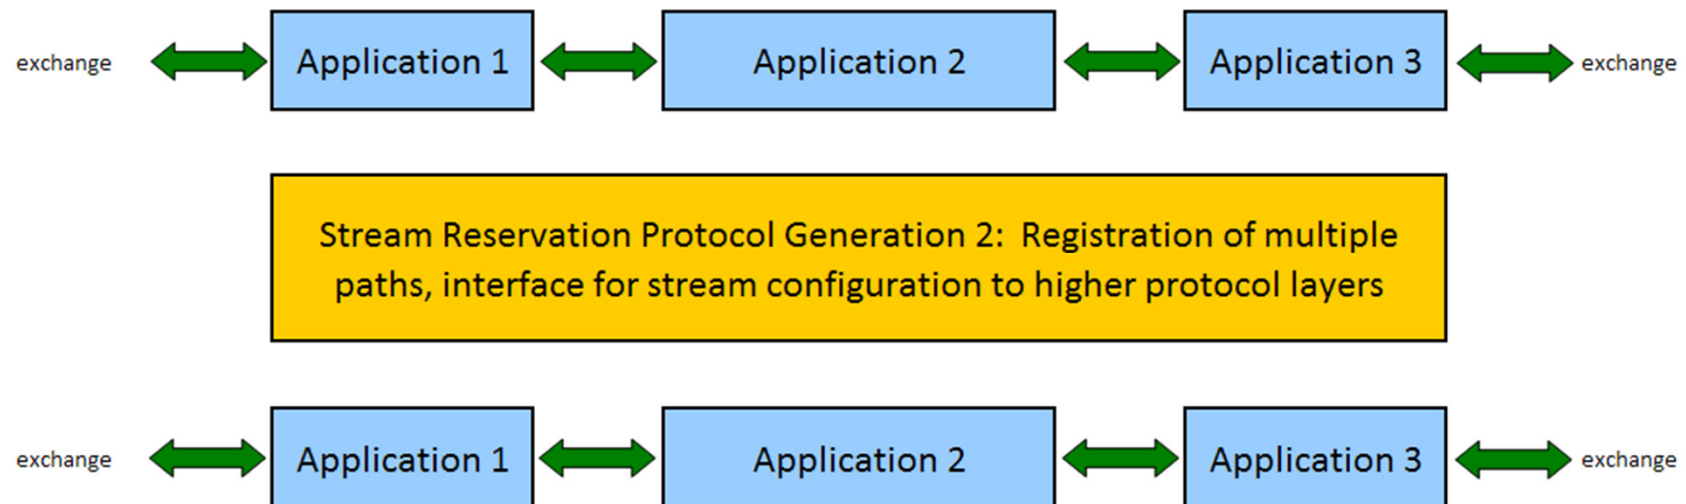
Assumptions and Requirements

Assumptions for the use of AVB Gen.2 in industrial and automotive environments:

- Standardization of Ethernet technology is still important (probably now more than ever), but new and very specific application demands will still outpace standardization or will never be truly standardized (in the near future)
- Especially with mission-critical applications, fault-tolerant design is paramount, it will be a basic requirement of most future Ethernet-based communication systems
- SRP needs to be able to adapt flexibly to new application requirements because it is instrumental in enabling fault-tolerant low-latency network design and needs to be able to work with today's and future Ethernet systems
- As applications come (and go), requirements and technologies on “higher” and “lower” layers of the overall architecture change. SRP must allow for those changes to happen.

Assumptions and Requirements

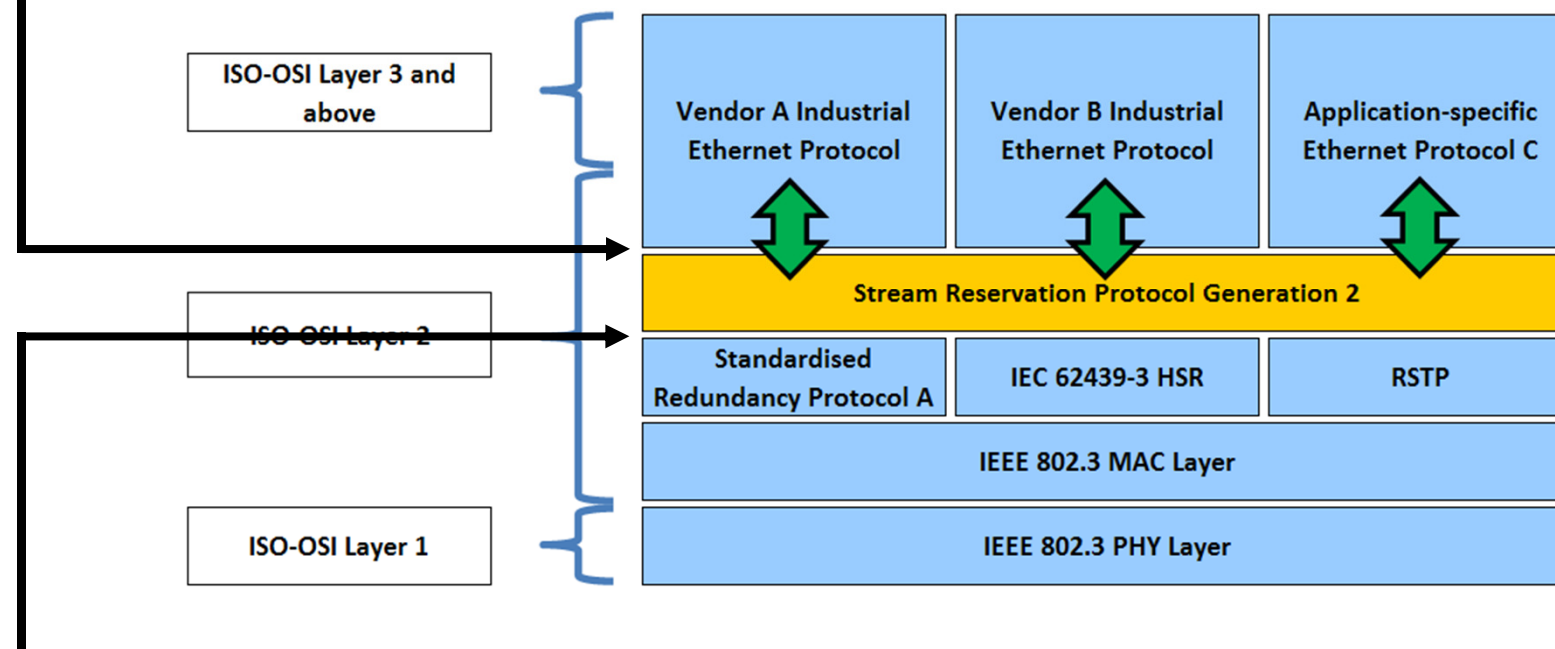
- Layering: To allow SRP to work with new (probably application-specific) technologies, the most important aspect of SRP Gen.2 is to allow those technologies to “attach” themselves loosely “on top” or “bottom up”



- SRP main extension proposal:
 - Registration of streams through all available paths (abstract from red. Control protocol and ensure a pre-calculable reconfiguration time
 - Service interface to higher layers to allow control of stream registration and transmission, (worst case) latency surveillance, etc...

Assumptions and Requirements

- SRP needs to offer an interface to higher layers to enable stream arbitration and control. → Interfacing to the application



- SRP needs to be able to operate with arbitrary (physical) topologies. These topologies are dependent on the redundancy control protocol, e.g. RSTP
→ Abstraction from the redundancy control protocol

Details can be found in: <http://www.ieee802.org/1/files/public/docs2011/at-kleineberg-AVB-media-redundancy-1111-v02.pdf>

Overview of the Proposal

Overview of the Operation

Overview of the Proposal

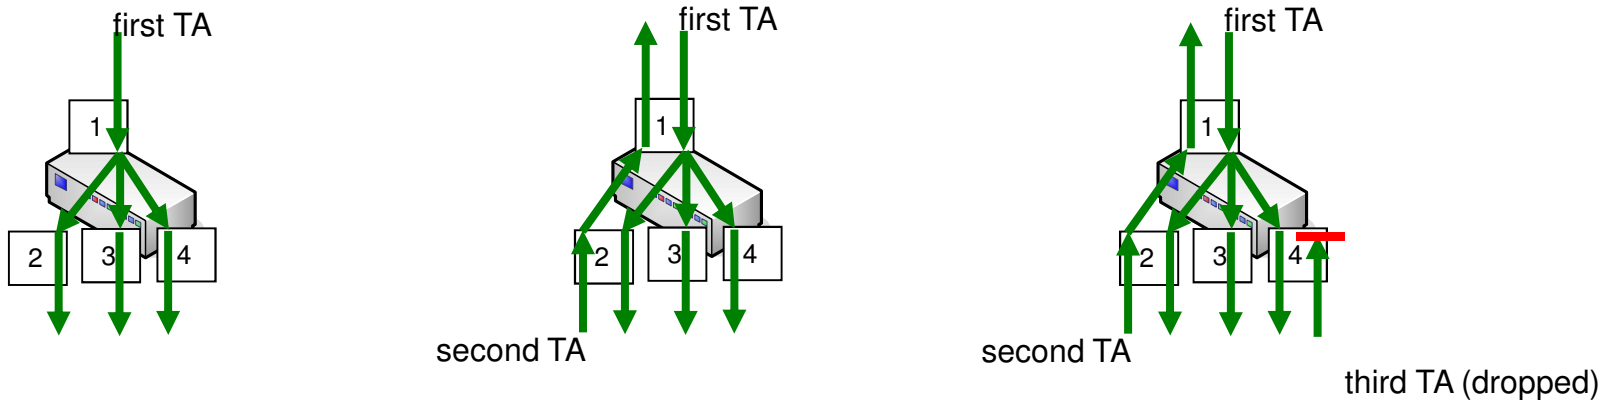
Proposal: „Mark“ streams that are meant to be sent redundantly and let bridges handle them accordingly

- Streams that are intended to be sent redundantly can be identified by a „redundancy identifier“ (to be defined, could be e.g. an attribute declaration) → Bridges track redundant streams by their ID and the redundancy identifier
- This „redundancy identifier“ can be either set by talkers that want a redundant network structure to handle its stream redundantly (or that have redundant network interfaces themselves) or it can be set by a bridge (e.g. a bridge that implements a redundancy protocol and that has a redundancy-unaware talker on one of its ports)
- Bridges produce (and consume) redundant stream registrations
- Talker Advertise and Listener Ready PDU's are transmitted and received through ports in „Discarding“ state. SRP does not follow the RSTP tree any more.
- Multicast stream frames still terminate at discarding ports, like any other „regular payload frame“ would.


Overview of the Proposal

Bridge Behaviour:

Talker Advertise




Listener Ready



Overview of the Proposal

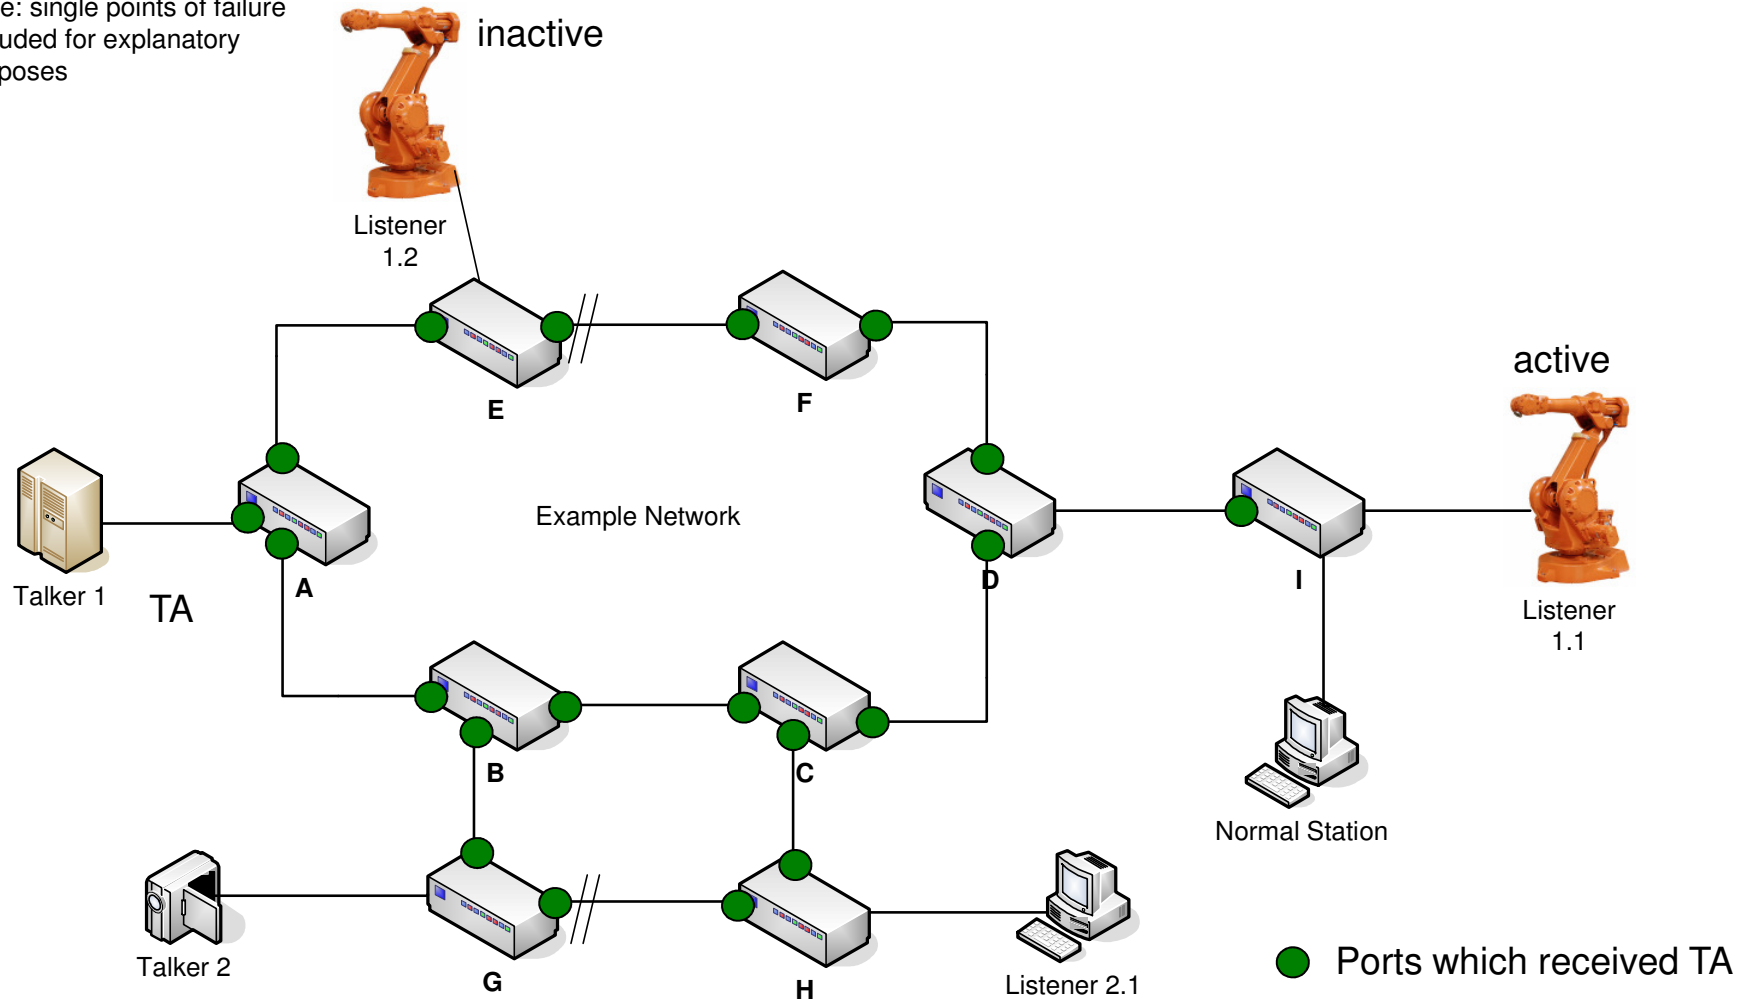
Note: single points of failure included for explanatory purposes



Talker advertisements are sent by bridges on all ports except the ports the same advertisement has been sent to already (and on which the same TA has been received)

Overview of the Proposal

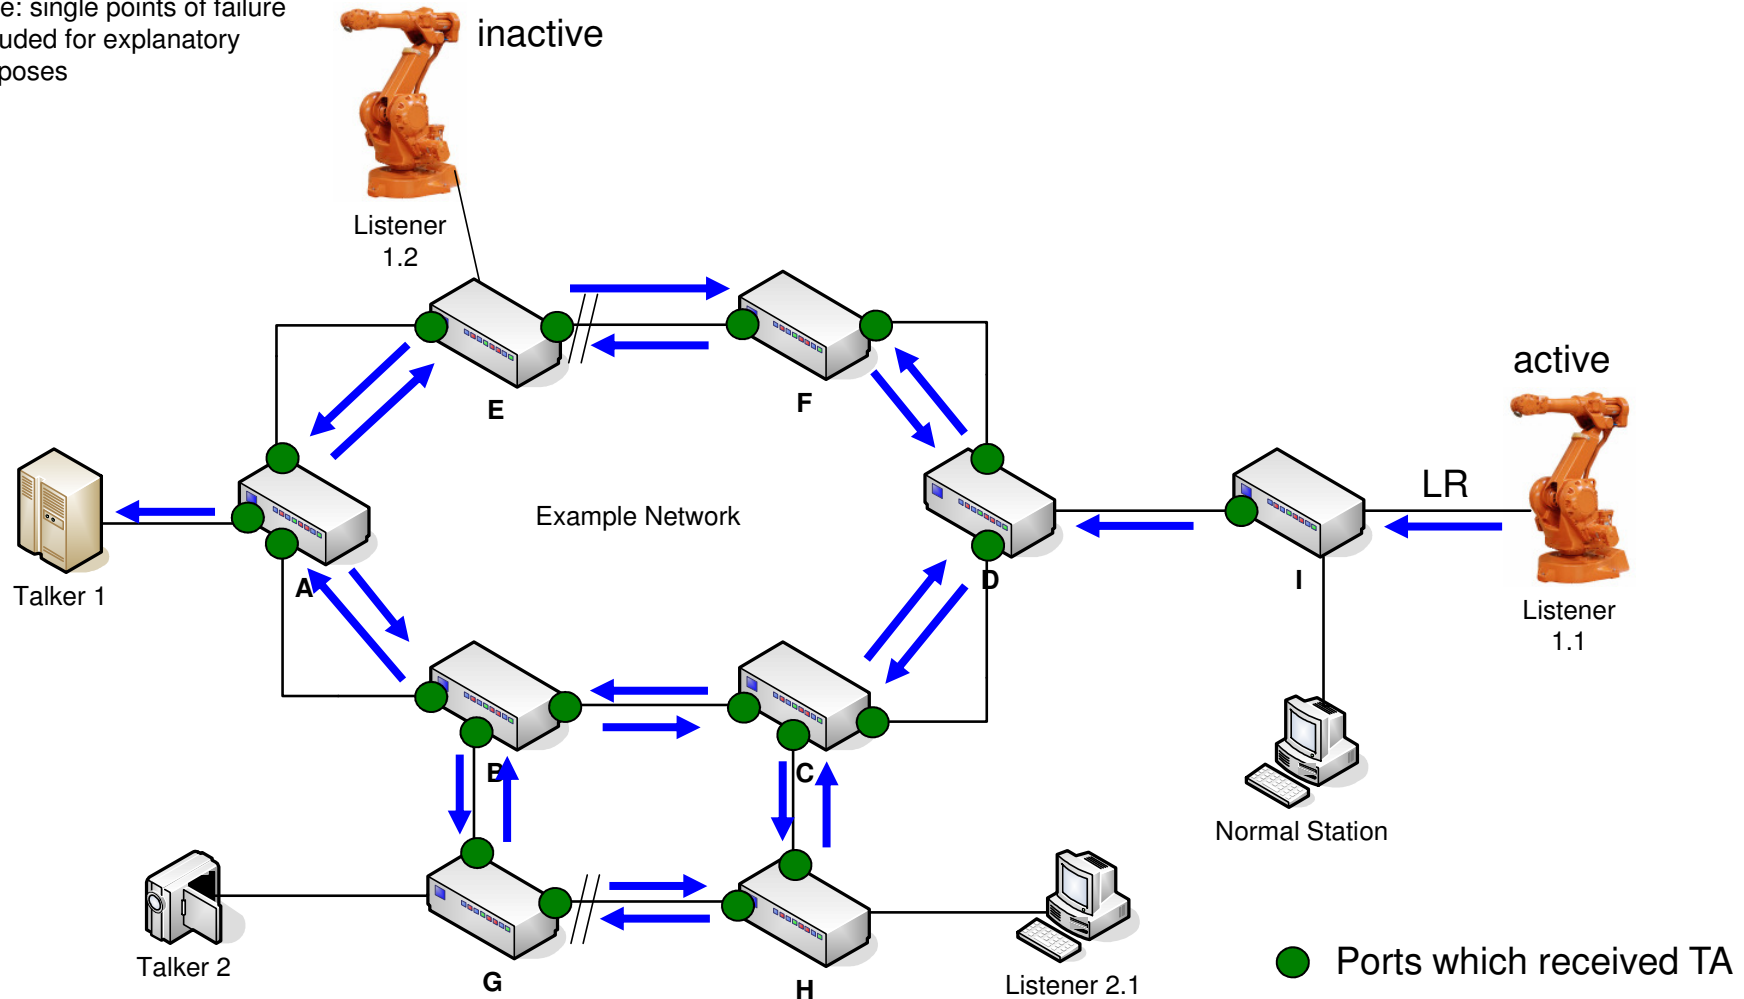
Note: single points of failure included for explanatory purposes



Bridges now know on which ports they can „reach“ the talker (The ports on which they recieved the TA)

Overview of the Proposal

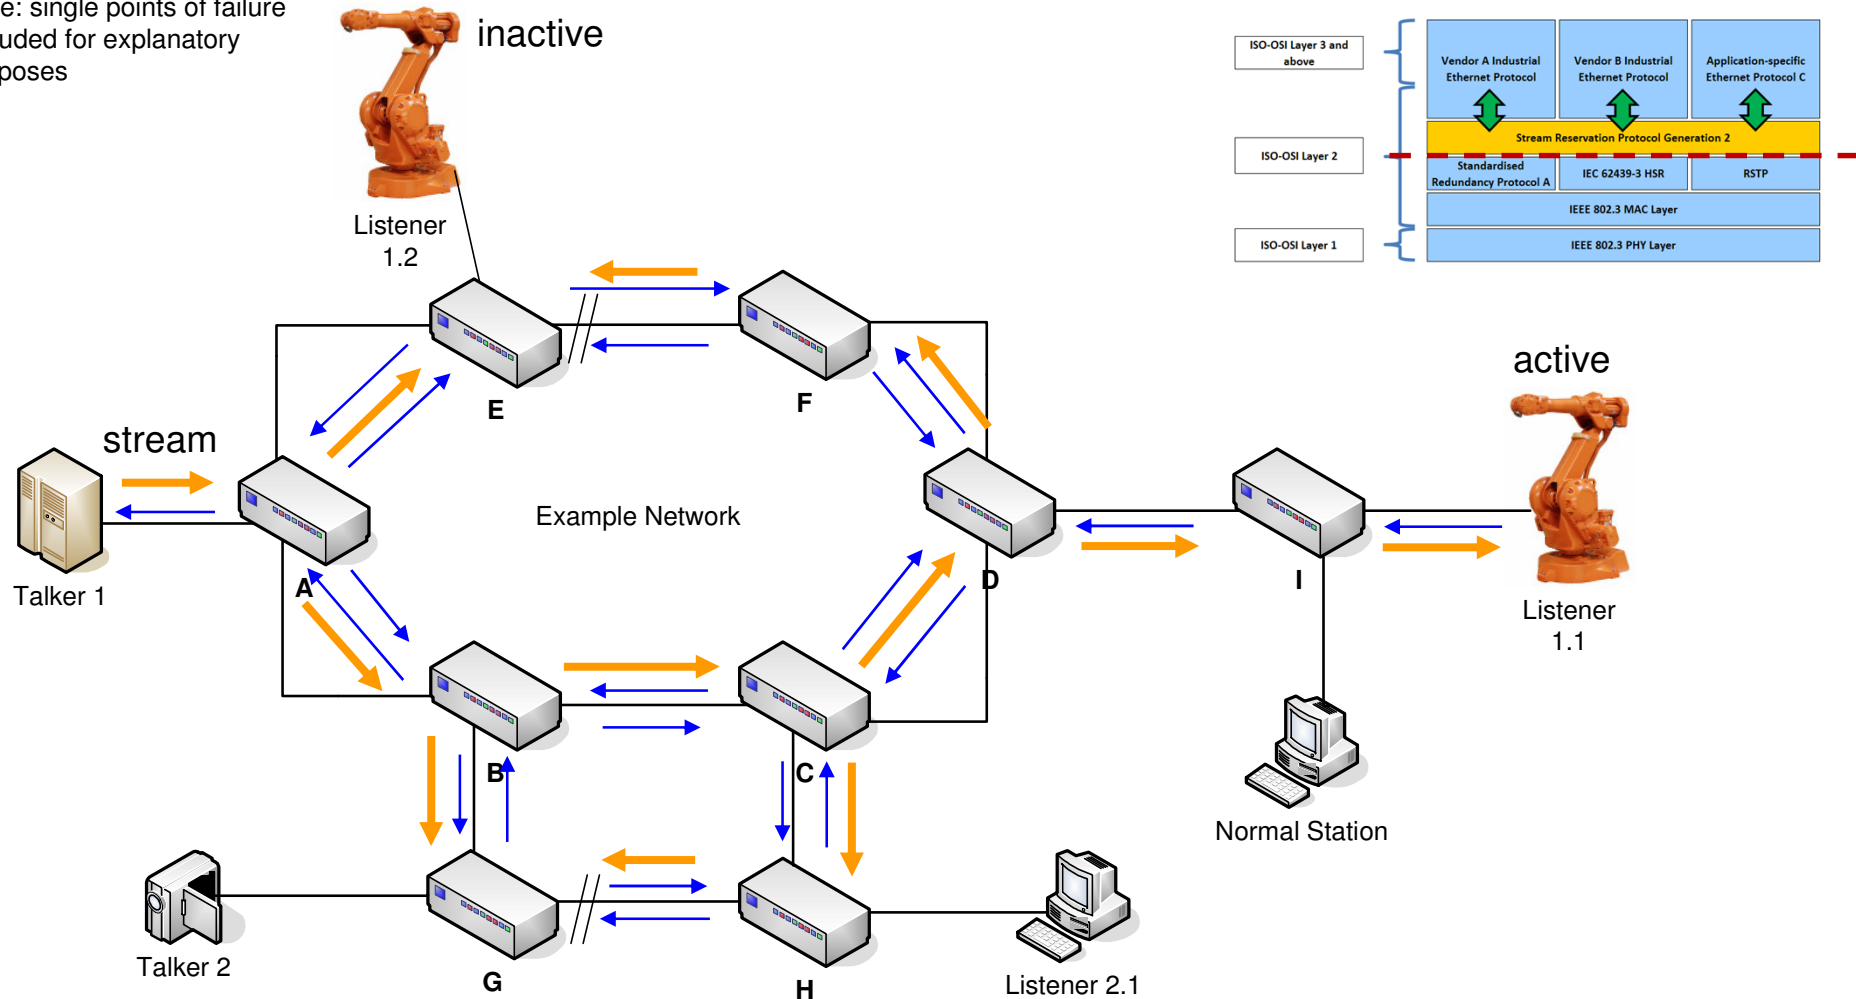
Note: single points of failure included for explanatory purposes



Bridges send „listener ready“ on all ports they received a „talker advertise“ (except the receiving port) – loops are prevented by the bridges keeping track where LR's were already forwarded.

Overview of the Proposal

Note: single points of failure included for explanatory purposes



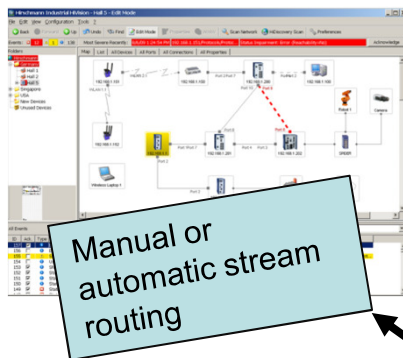
This effectively **decouples** Stream registration and transmission from the redundancy control protocol.

It also makes sure that $T_{rec_overall} = T_{rec_steam} + T_{rec_redundancy}$ can be precisely pre-determined.


Overview of the Proposal

- The “birds eye” view: Different use cases


Ind. Automation SCADA
with Topology Knowledge



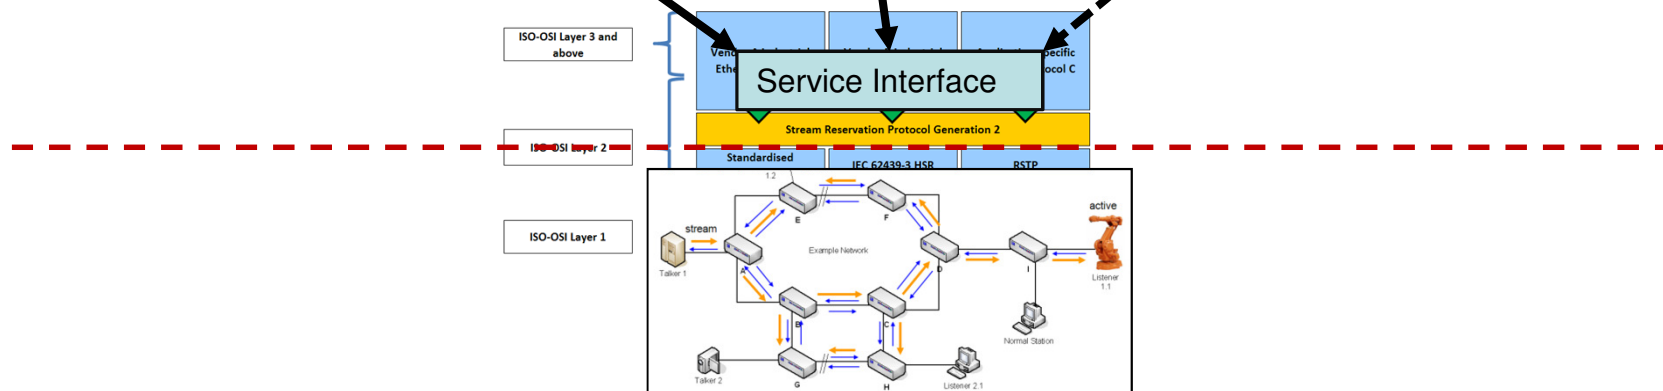
Automotive fixed network
configuration



Home or IT Network




Configuration / Information Exchange



Example application scenarios and benefits

Example application scenarios and benefits

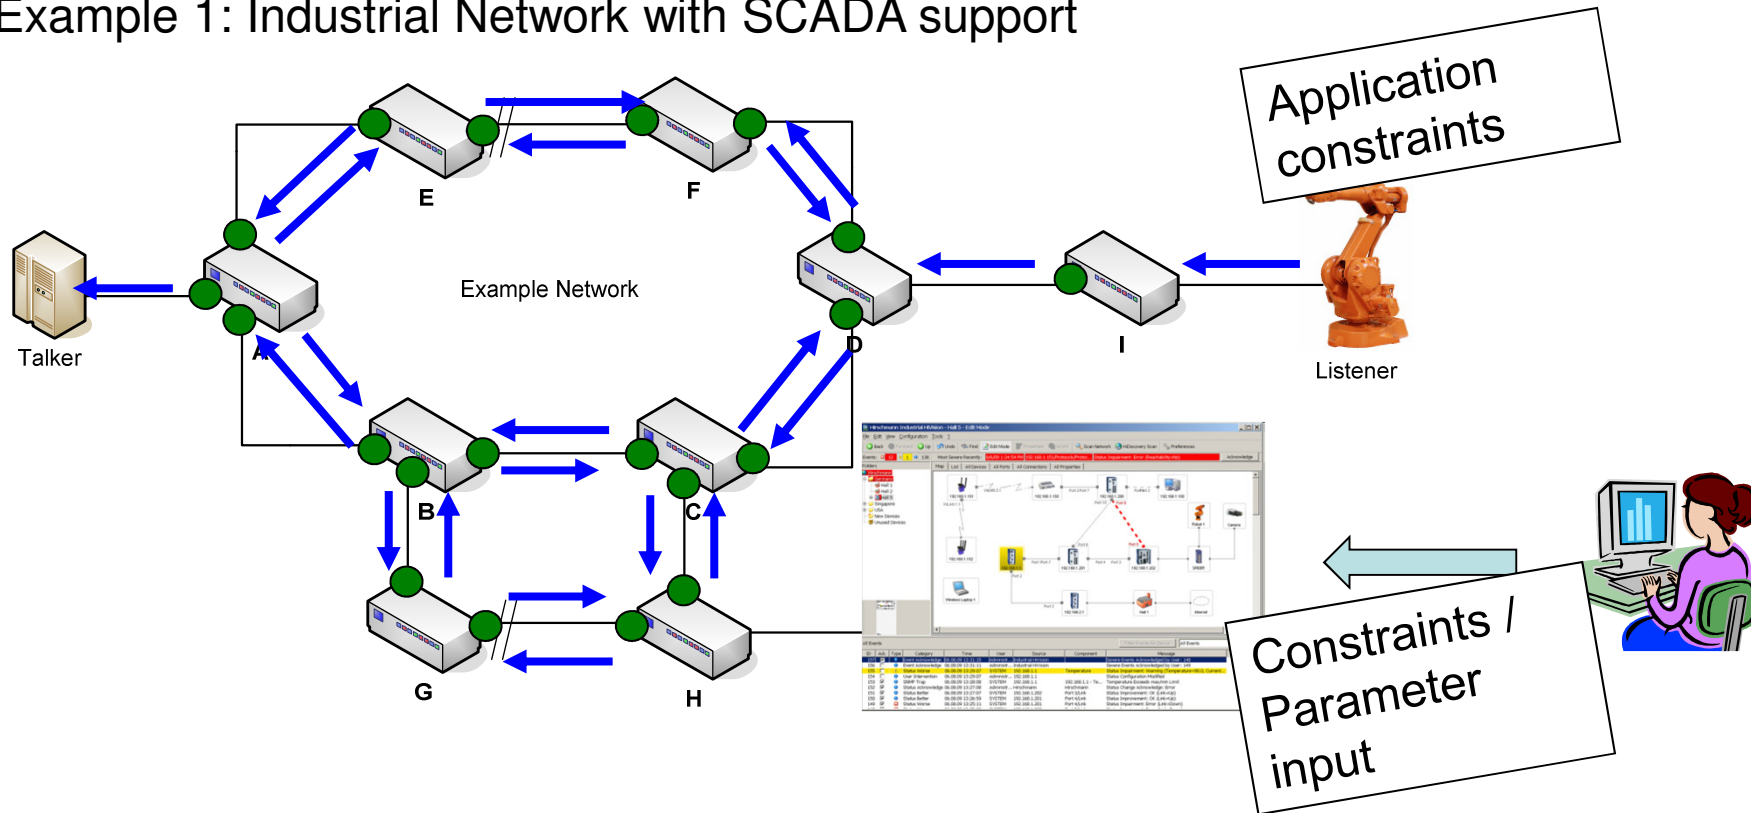
Example 1: Industrial Network with SCADA support



- RSTP used for Redundancy
- All paths are pre-registered by MSRP
- SCADA extracts from Network: Available Paths for Streams, Topology/discarding ports and accumulated/projected latency

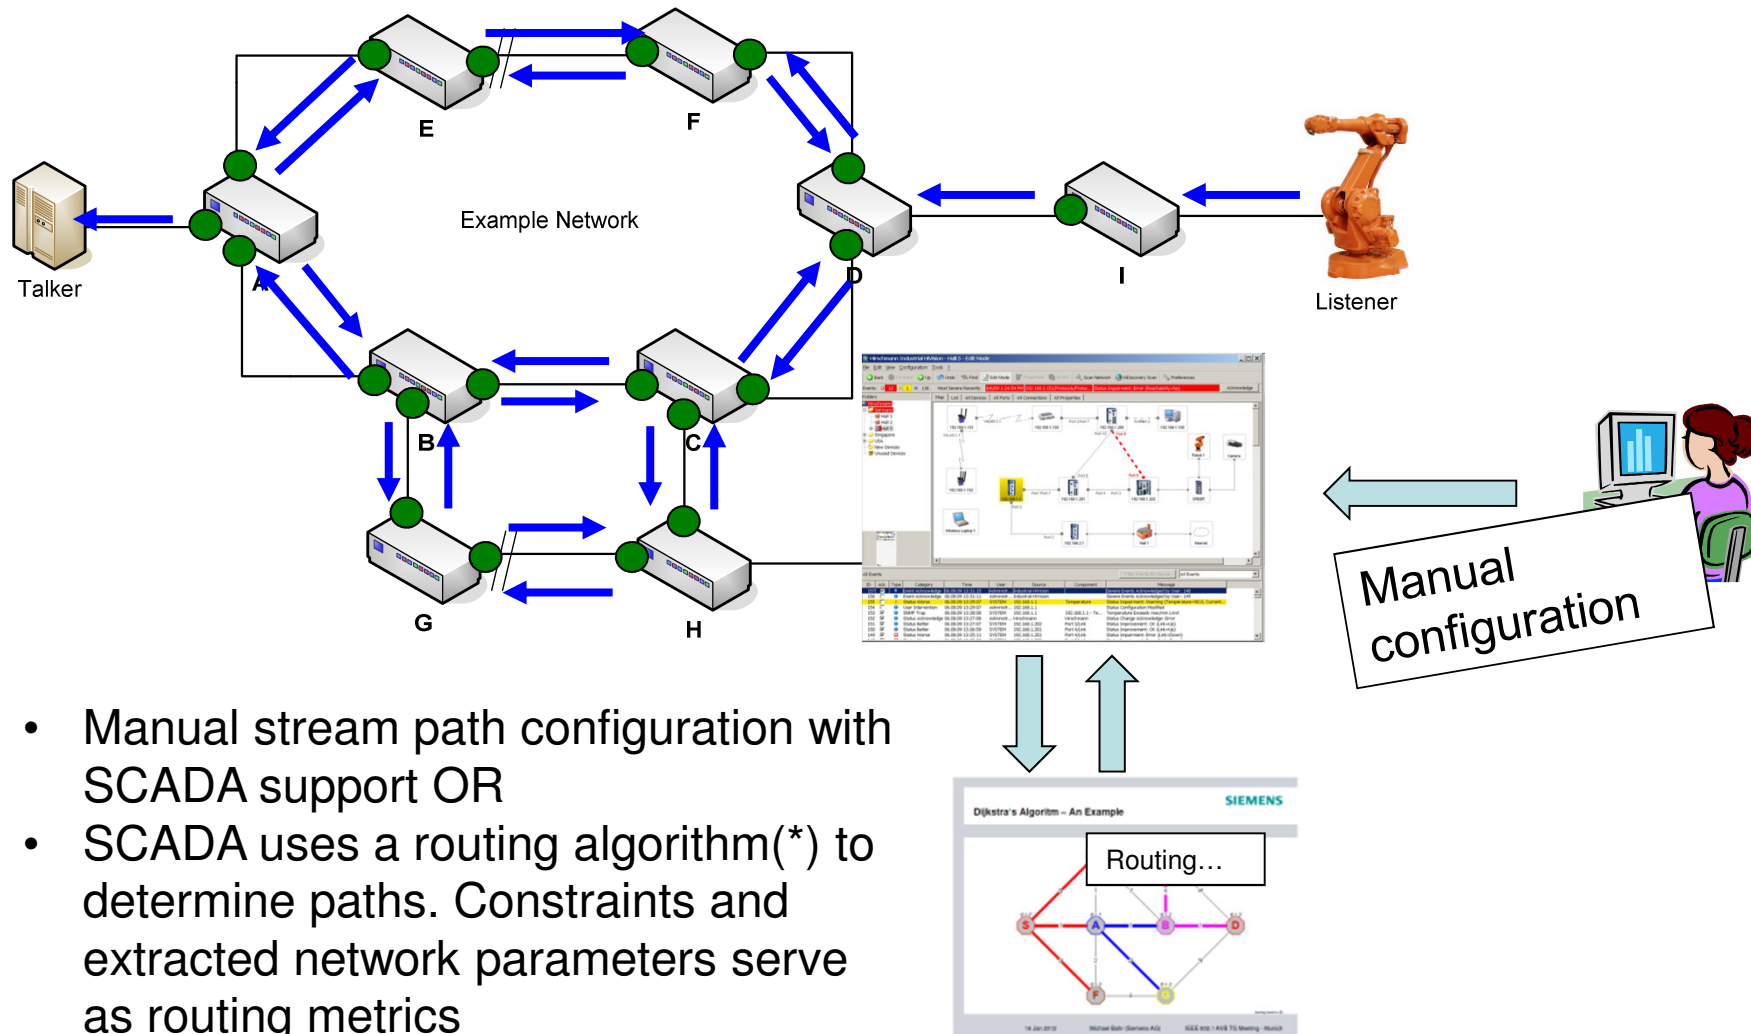
Example application scenarios and benefits

Example 1: Industrial Network with SCADA support




- Engineering input: constraints / parameters concerning e.g. availability, latency OR
- Automatic discovery through constraints definitions on the applications

Example 1: Industrial Network with SCADA support

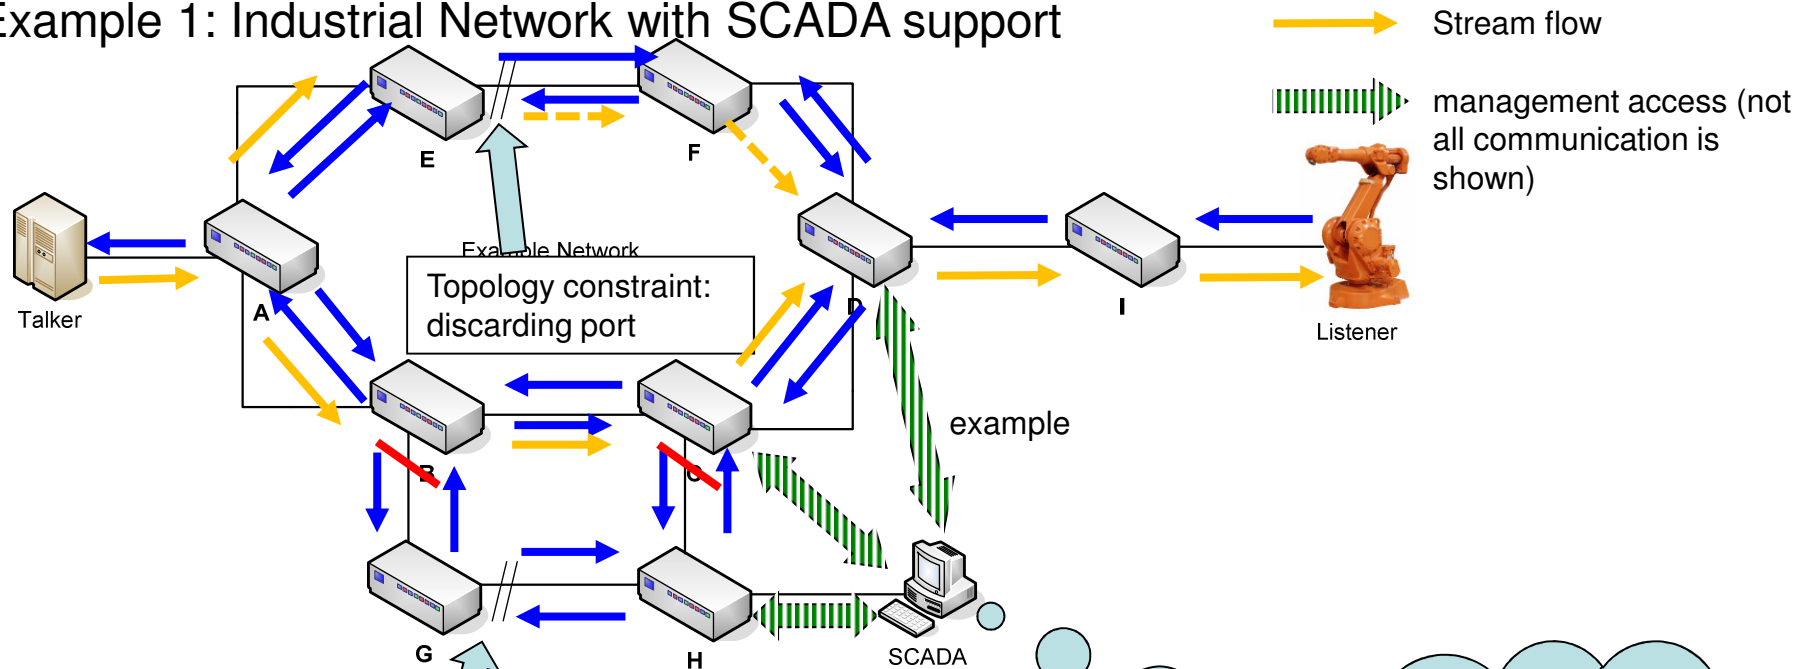


Example 1: Industrial Network with SCADA support



Example application scenarios and benefits

Example 1: Industrial Network with SCADA support



The subring with G and H gets removed from the Topology to improve worst case E2E latency

Constraints:

Need min 1 active path and one path for fault tolerance

Max Latency for Listener = x

→ Calculate logical stream topology...

Example 1: Industrial Network with SCADA support

Integration of Shortest Path Bridging - SPB:

Option 1: „SPB on top of Multi-Path MSRP“


- + Application constraints can directly influence available paths for SPB to make sure requirements are met
- + SPB can be omitted where it is not feasible/needed (see automotive)
- Additional „flooding“ of MSRP telegrams to the whole network (but no multiple loops)

Option 2: „Multi-Path MSRP on top of SPB“

- + No „chatty“ MSRP telegram „flooding“ through the network
- + Application still can „see“ relevant parameters e.g. worst case latency
- Application cannot directly influence available paths to configure stream flow according to requirements
- SPB will be an essential part for using Gen.2 MSRP

Example application scenarios and benefits


Example 2: Automotive backbone network: reduced to the essentials



- Network and communication relations are fixed at engineering time...
- ... and installed in NV RAM at vehicle production time
- Topology is fixed, latency and device/application parameters are known
- Instead of a SCADA system configuring stream flow...
- ... the fixed configuration is read from NV RAM and is used to configure the network through the service interface on vehicle start


Example application scenarios and benefits

Example 3: Open customer premise network with audio/video service




Can work completely unattended or can be (partially) engineered– redundant paths improve service quality for customers

Example 3: Open customer premise network with audio/video service



In case of communication interruption...

Example 3: Open customer premise network with audio/video service



...after RSTP reconfiguration, a backup path takes over „seamlessly“

For a future SRP Gen. 2 that goes beyond the scope of Audio and Video applications, the following things should be considered:

- Redundancy for fault-tolerance is a mandatory requirement for use in mission-critical applications...
- ...but it can also be beneficial for “standard” applications
- Total network reconfiguration time needs to be pre-determinable
- The overall system design should not be fixed to a specific (redundancy) protocol (AVB Gen.1 – RSTP should not be replaced with AVB Gen.2 – SPB)
- ... because this will limit the technology scope to areas where SPB implementation is feasible

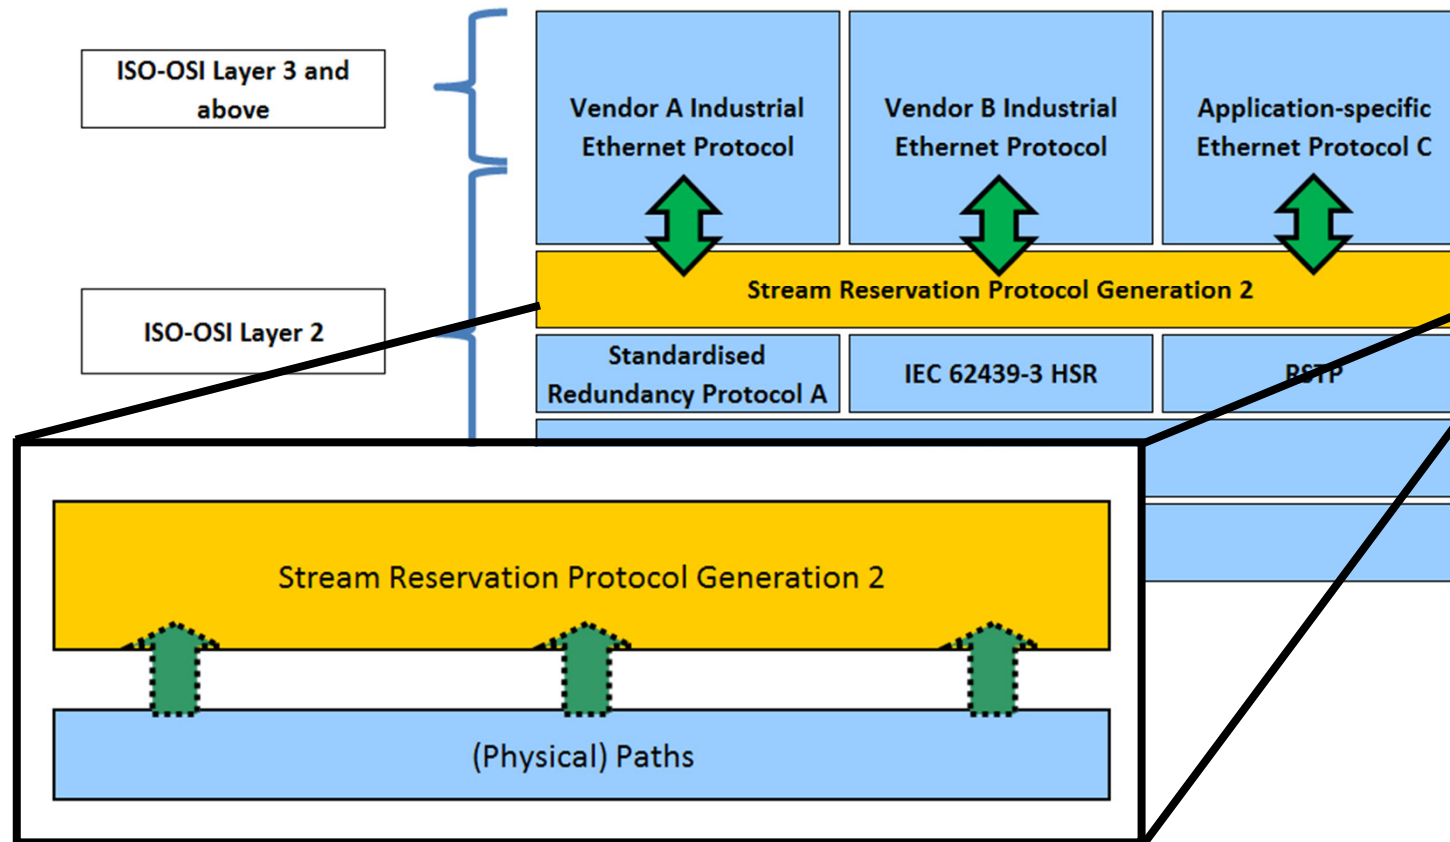
To make AVB Gen.2 a success, make sure it can be used in the target markets!

Thank you for your attention!

Questions?


Backup Slides

Future-proof SRP: „Layering“ and interfaces to other technologies



- Physical (and Logical) Topology are imposed on SRP
- SRP Gen.1 still follows the RSTP logical topology
- SRP Gen.2 observes and registers streams on all available paths, ignoring discarding ports for stream registration

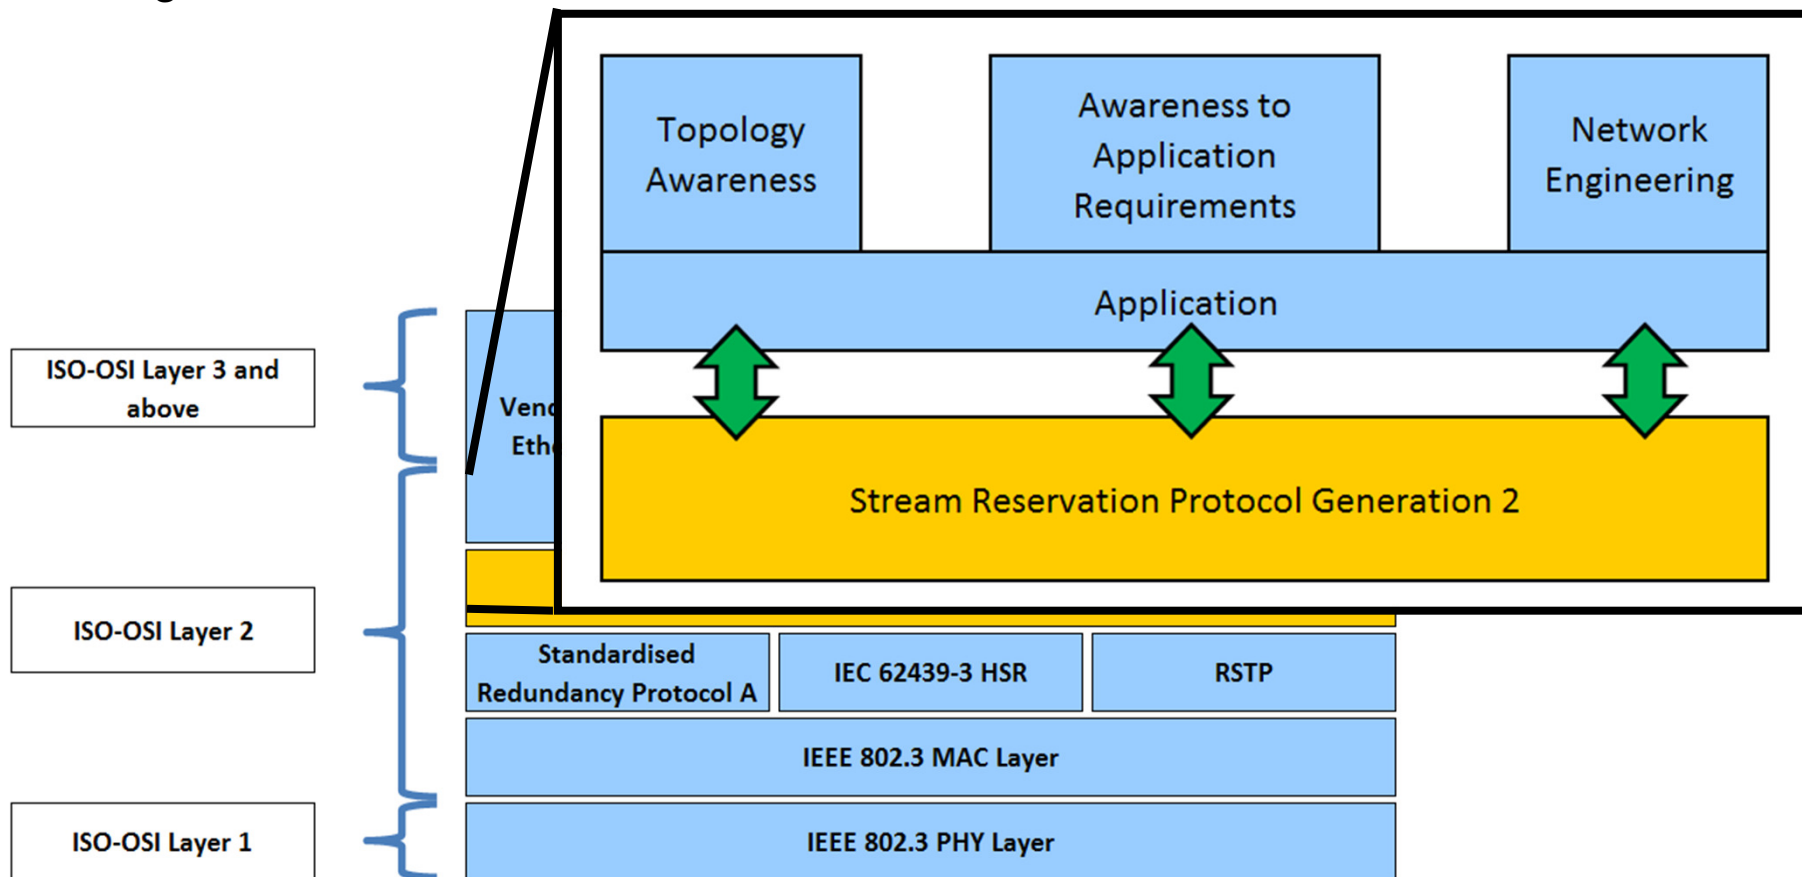
Future-proof SRP: „Layering“ and interfaces to other technologies



- For Gen.2 registration of multiple paths, see [slides_singapore]
- This allows SRP to achieve the switchover times that are in line with the underlying redundancy control protocol. (e.g. IEC-HSR, RSTP, SPB,...)
- The used redundancy protocol depends on application requirements

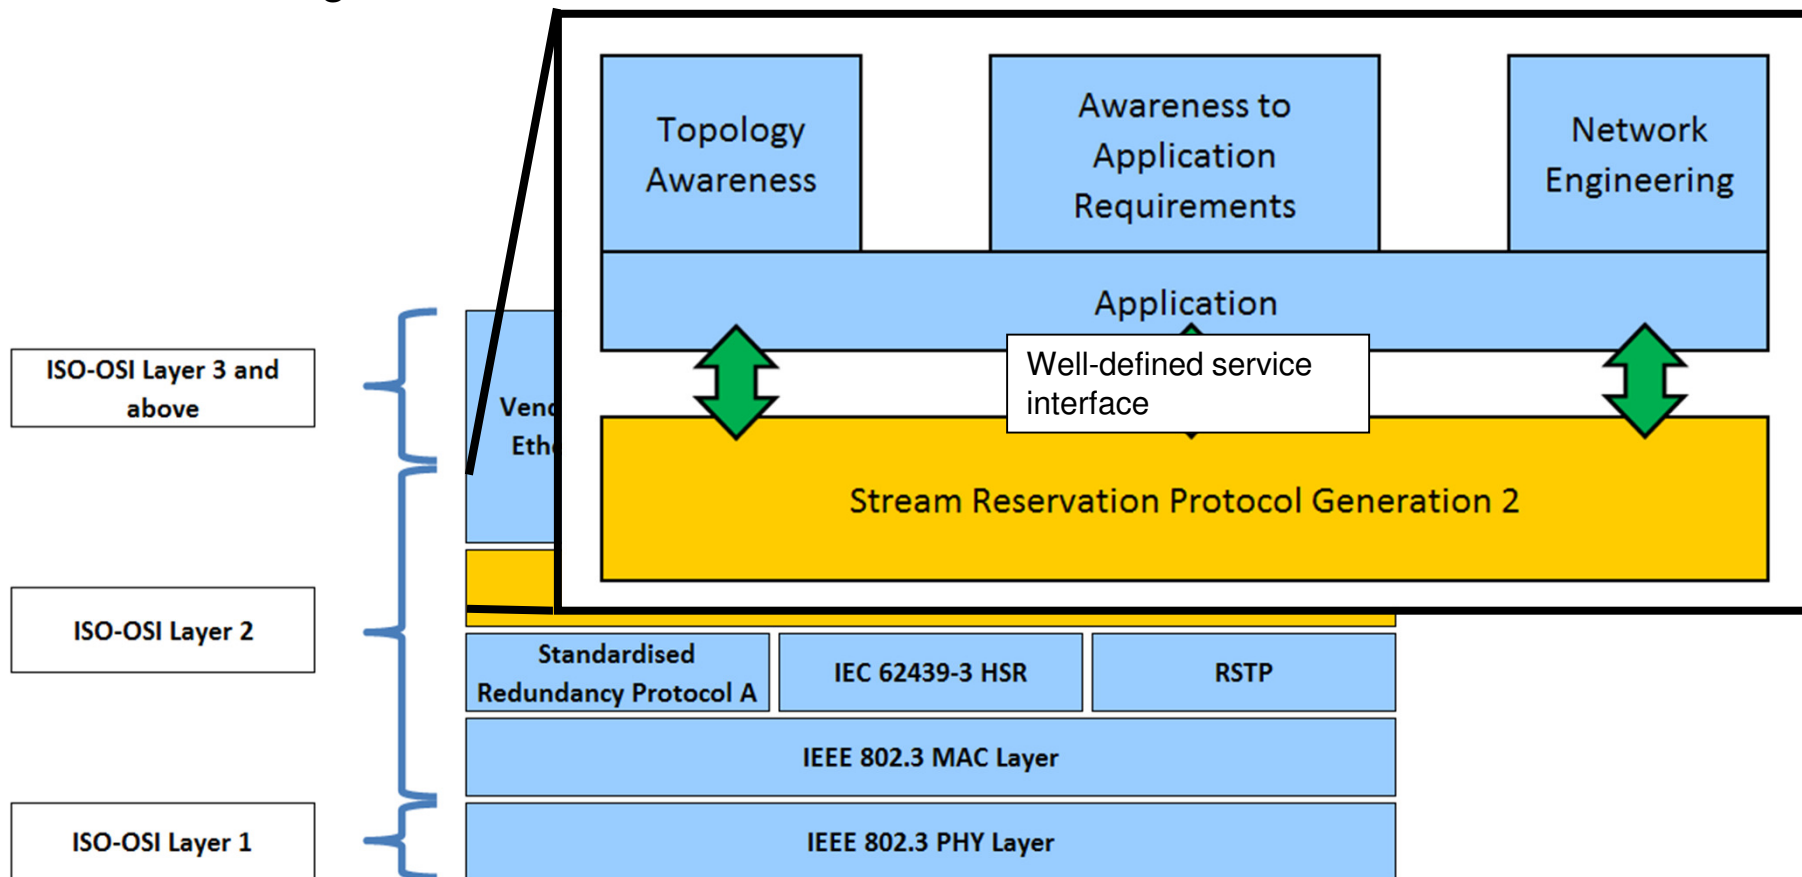
Future-proof SRP: „Layering“ and interfaces to other technologies

- Higher Layer entities usually have a complete topology awareness (e.g. Industrial Engineering Tools, SCADA systems, ...)
- Topology awareness and application req. awareness are used to configure / engineer stream flows




Future-proof SRP: „Layering“ and interfaces to other technologies

- Higher Layer entities can: enable or disable streams entirely, control stream flow through enabling/disabling bridge ports, etc...
- Higher Layer entities are provided with information on streams and configure SRP through a well-defined service interface




Future-proof SRP: „Layering“ and interfaces to other technologies

- Information from SRP: e.g. Maximum worst case latency from talker to listener, based on multiple paths (i.e. all latency information for all paths registered)
- Each SRP Gen.2 device must provide worst case latency information independently and the worst case must observe all paths from talker to device




Future-proof SRP: „Layering“ and interfaces to other technologies

- Higher layer entities could be (also see [slides_sanfrancisco]):
 - Not present at all (in that case, streams on all paths will be registered)
 - Automated or non-automated network engineering (e.g. Industrial Ethernet Engineering tool with algorithmic support)
 - Fixed configuration (for 100% static network configurations e.g. automotive networks)
 - ...




Example application



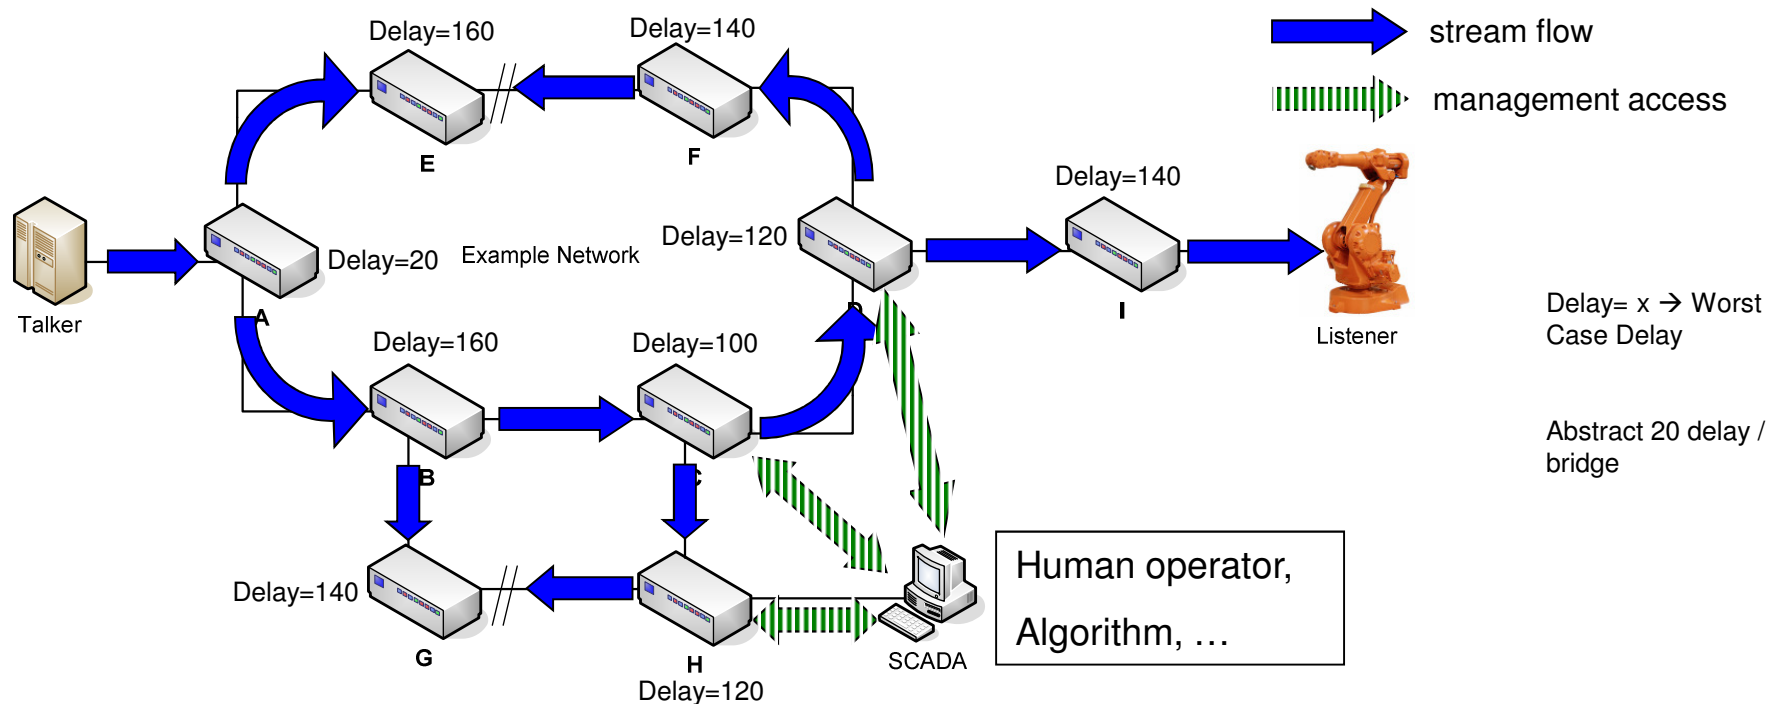
- A single talker and listener (Industrial control application) want to communicate through a fault-tolerant network
- The redundant paths in the network are administrated by RSTP
- A SCADA system is in place as an engineering workstation. It has full topology knowledge (e.g. through SNMP and LLDP) and management access to all bridges in the network (e.g. through SNMP)

Example application




1. SRP registers the redundant stream on all available paths from talker to listener (details on how this might be done: [slides_singapore])
2. The SCADA system collects all stream data from all the bridges, e.g. stream data, latency... (Only exemplary access is shown above to not overburden the picture)

Example application



3. The SCADA system displays the topology, together with the stream flows and the worst case latency (when more than one path is available)
4. From this information, a human operator through a network engineering tool, an algorithm with application-specific knowledge, etc... can influence which paths are to be configured for stream transmission

Example application



5. In this case, a human operator decides to cut off the sub-ring through G and H from this stream to reduce worst case latency, at the cost of some fault-tolerance
6. In another use case (and with other requirements in the background), the outcome of this decision could have been different!... E.g. when the additional fault-tolerance outweighs the improvements to latency