

25G PHY TYPES

Jeff Slavick, Avago

Adee Ran, Intel

Goals

- Move towards a decision about number of PHY types
- Satisfy the project objectives:
 - Define a single-lane 25 Gb/s PHY for operation over links consistent with copper twin axial cables, with lengths up to at least 3m
 - Define a single-lane 25 Gb/s PHY for operation over links consistent with copper twin axial cables, with lengths up to at least 5m
- Satisfy the CSD, especially:
 - Broad market potential
 - Broad sets of applicability
 - Economic Feasibility
 - Consideration of installation costs
 - Consideration of operational costs (e.g., power consumption)
- Address other apparent needs:
 - Low-latency operation
 - Interoperability

Use cases

- **Engineered network**
 - Primary optimization factors: cost, power, latency...
 - Supported and desired modes are known in advance and pre-configured.
 - Interoperability is guaranteed by design.
- **Plug-and-play**
 - Primary optimization factors: install/service time, user experience.
 - Users may care little about selecting options and configuring modes – “make it work”!
 - Interoperability is assumed by compliance with a standard.

Modes, Optional, Mandatory?

- *Mode*: one out of several behaviors allowed by the standard
- *Optional mode*: a mode that may or may not be implemented by a compliant device
 - Example: clause 74 Base-R FEC is optional for most 10G PHYs.
- *Mandatory mode*: a mode that must be implemented by a compliant device
 - A mandatory mode may not be used in some cases, another mode (optional or mandatory) can be used instead.
 - Example: operation without clause 74 is mandatory; it must be selected if partner does not support clause 74.
- Some modes have to be used on both ends to interoperate
 - This can be done by auto-negotiation, e.g. usage of clause 74 in current PHYs.
 - Can also be done using management, e.g. “disable training”.

Auto-negotiation

- **Clause 73 defines information passed and priority-based resolution**
 - HCD determines PHY choice.
 - FEC ability/request bits determine usage of FEC.
 - EEE ability bits determine usage of LPI.
- **Management can control which supported abilities to advertise**
 - This is out of scope of clause 73 (and 802.3 in general).
 - Partial advertisement typically *reduces* interoperability.
- **Link should be resolved in a single pass if abilities match**
 - This is the spirit of clause 73.
 - This is what users expect as normal behavior.
 - Multi-pass management algorithms might not be interoperable.

FEC

- The encoding “menu” includes
 - **RS-FEC**: required to meet the 5 m objective, high latency, largest gate count
 - **Base-R FEC**: is required to meet the 3 m objective, medium latency, low gate count
 - **No FEC**: Although not part of the project objectives, there is an apparent desire to enable it in adequate links. (See [andrewartha_3by_01a_0115](#))
- Decoding and encoding must match... AN can be used for that
 - What should be advertised?
 - What are the resolution rules?
- Which mode is mandatory?

Possible paths for single PHY

Single PHY type (CR)

This choice must be made by the task force!

One mode

Low latency or 5 m reach?
Impacts Broad Market Potential
Unsuitable for engineered networks

Multiple modes,
only Base-R FEC
mandatory

5 m reach objective not
guaranteed (depends on an
optional feature)
Unconventional, might
confuse the market

Multiple modes,
only RS-FEC
mandatory

Over-design
Impacts Economic
Feasibility
Unsuitable for
engineered networks

Multiple modes,
no mandatory
FEC

No guaranteed common
capability even for 3 m
Impacts interoperability
Unsuitable for plug-and-play

Possible paths for multiple PHYs

Multiple PHYs (CR-L and CR-S)
Advertise capabilities through AN

This choice can be made by vendors!

Single capability device –
only CR-S

Optimized for low latency
No over-design

Can be used for networks using
only CA-S cable assemblies
(engineered)

Single capability device –
only CR-L

Optimized for maximum reach and
support of other standards
Can be used with dual-capability
device over any cable

Idea can easily be
extended to three
PHY types

Dual capability device –
both CR-S and CR-L

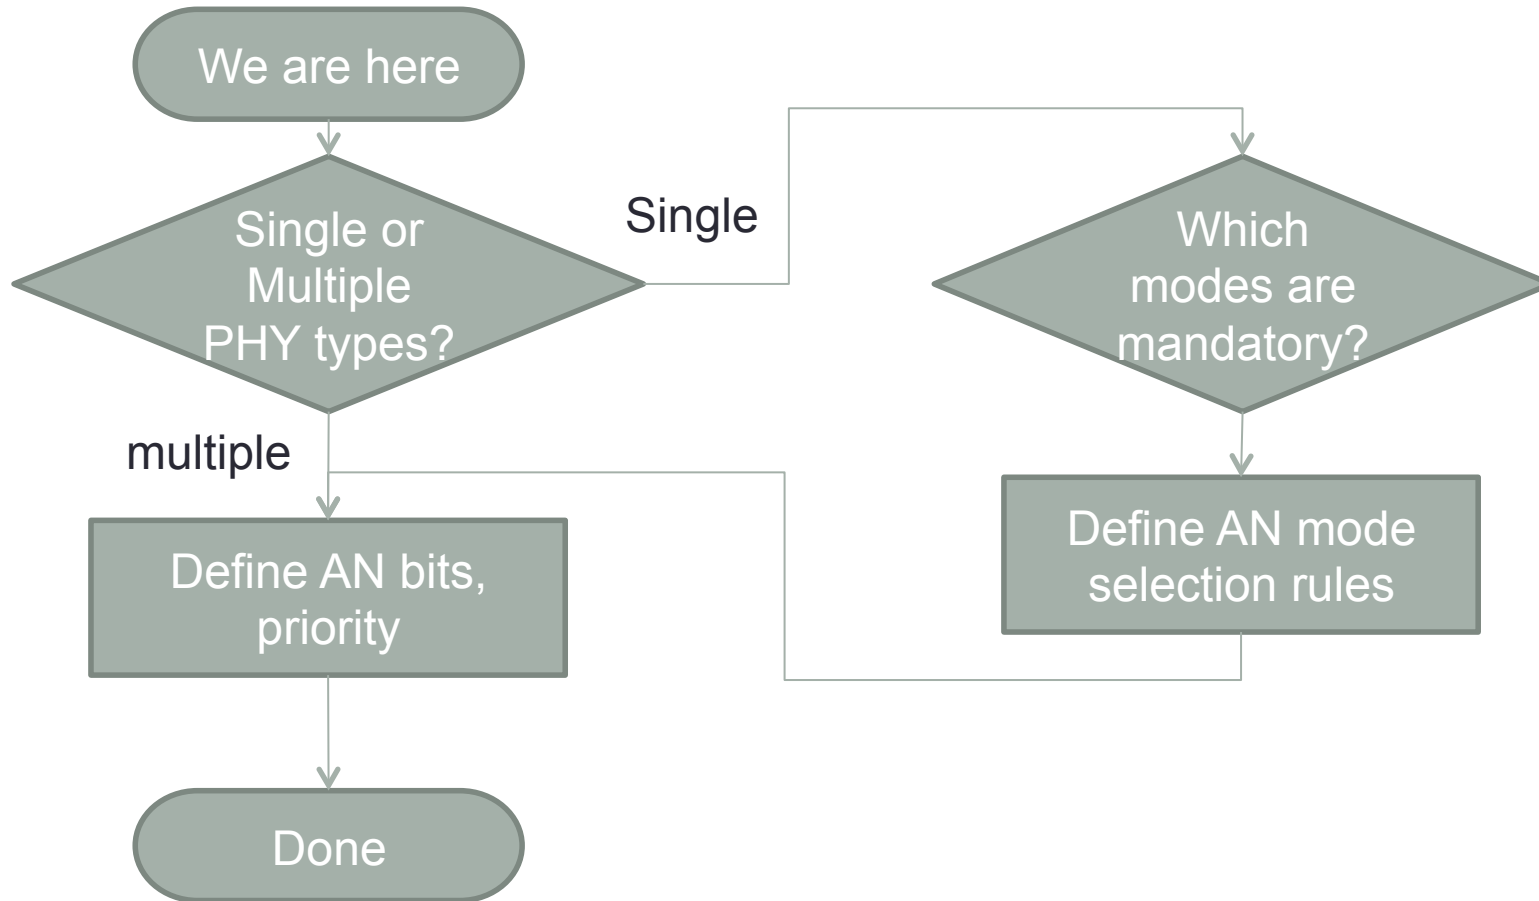
Optimized for interoperability
Can be used with any single-
capability device over compliant
cables for the common operating
mode

May have higher cost

AN resolution

- If we go towards several PHY types, which can be advertised separately through AN – we should address HCD resolution.
 - If only one capability is common to both partners – the resolution is clear.
 - If more than one capability is common, one should be selected based on priority.
 - Does it matter which one has higher priority?
 - Current clause 73 has predefined priorities for each PHY – not user selectable. This prevents deadlock situations.
 - What if we set a priority order, but a user has different preferences?
 - Possible solution: disable advertisement of undesired modes
 - Other solutions can involve reading media type and advertising accordingly.
- If we go to one PHY type, HCD is not an issue
 - If several FEC options are available, a FEC resolution is still required.
 - Simple rule in 73.6.5 may be sufficient, or we may need new logic.

Decision making flowchart



PCS, FEC area cost and performance

	Gates	% of total	35 dB BP	3m Cable	5m Cable
Clause 49	45k	9%	No way	Possibly	No way
Clause 74	80k	15%	Doubtful	Likely	Doubtful
Clause 108	400k	76%	Likely	No problem	Likely
Total	525k				

	Area	PCS % of PHY
PMD/PMD	X	
CI 49	X * 0.05	4.5%
CI 49, 74	X * 0.13	11.7%
CI 49, 74, 108	X * 0.65	34.8%