IEEE P802.3bt Layer 2 Ad Hoc

Wael William Diab – Huawei
(on behalf of the Ad Hoc)

May 2015

Agenda

- Summary of Calls
- Existing Features
- New Features
- Next Steps

Summary of Calls

Summary of Calls

- Conducted 4 Calls
- Agreed to framework for driving layer 2 work
  - Define layer 2 feature in one slide summary
    - Multiple discussions to distill out clear definition and avoid confusion down the process
  - Identify impact to layer 2 mechanism: TLV only OR TLV + SM changes
  - Generate TLV and baseline text
- Reviewed feature list from straw polls
  - Consensus to pursue all features. Drop Enhanced MPS
  - Assignment of leads for each area
    - Autoclass – Lennart. PD Watchdog – Lennart. PSE Measurements – Yan. PD Measurements – Yan. Reserved class and types – Sesha
- Entertained new features proposed by attendees

Summary of Attendance

Call 1

- Chad Jones
- Bruce Nordman
- Bryan Moffitt
- Fred Schindler
- Gaoling Zou
- John Skinner
- Pete Scruton
- Sesha Panguluri
- Yan Zhuang
- John Wilson
- Wael William Diab

Call 2

- Chad Jones
- Bruce Nordman
- Fred Schindler
- Gaoling Zou
- John Skinner
- Chris Bullock
- Dylan Walker
- Lennart Yseboodt
- Yan Zhuang
- David Tremblay
- Miklos Lukacs
- Ken Bennett
- Wael William Diab

Summary of Attendance

Call 3

- Yan Zhuang
- Fred Schindler
- Chad Jones
- John Skinner
- Miklos Lukacs
- Wael William Diab
- Chris Bullock
- Dylan Walker
- Bruce Nordman
- David Tremblay

Call 4

- Yan Zhuang
- Fred Schindler
- Chad Jones
- John Skinner
- Wael William Diab
- Dylan Walker
- Bruce Nordman
- Gaoling Zou
- Ken Bennett
- Pete Scruton

Existing Features

Overview

- Unanimous consensus to pursue all features
- Unanimous consensus to drop Enhanced MPS
- Progress on all assigned areas and consensus on 1-slide summaries
  - No work done on reserved class and types
- All targeting TLV changes and only

Autoclass

- Autoclass is a mechanism for optimal power classification, see
  - yseboodt_3_1114.pdf,
  - yseboodt_1_0115.pdf and
  - yseboodt_1_0115_baseline.pdf
- An Autoclass TLV would add the capability to tell the PSE to start an Autoclass procedure. The PSE would measure the actual PD consumption and is then allowed to reduce the power budget.

Power Cycle

- A PD can request the PSE to turn off the power. The PSE would then next start detection and classification again. This allows a PD to completely reboot / power cycle itself.

PSE Measurement

- Problem Statement
  - Actual losses in the channel maybe less than worst case scenario described by the standard. Unused budget can be communicated to the MDI partner and recouped / reused
  - Diagnostics of the installation and lifetime degradation of the cable etc.
- Feature Overview (optional feature)
  - Channel measurements are conducted by the PSE to determine actual channel loss vs. assumed worst case loss
  - Measurements communicated from the PSE to the PD via layer 2 mechanism to recoup/reassign the unused budged
- Feature Items for Discussion
  - Measurements of PSE voltage VPSE and port current IPORT-PSE to be exchanged.
  - Coordination with work on "Unused Power Owner"
- Changes
  - Additional TLV fields and associated management attributes / MIB amendments
  - No changes to the TLV processing state machines

PD Measurement

- Problem Statement
  - Actual losses in the channel maybe less than worst case scenario described by the standard. Unused budget can be communicated to the MDI partner and recouped / reused
  - Diagnostics of the installation and lifetime degradation of the cable etc.
- Feature Overview (optional feature)
  - Channel measurements are conducted by the PD to determine actual channel loss vs. assumed worst case loss
  - Measurements communicated from the PD to the PSE via layer 2 mechanism to recoup/reassign the unused budged
- Feature Items for Discussion
  - Measurements of PD voltage VPD and port current IPORT-PD to be exchanged.
  - Coordination with work on "Unused Power Owner"
- Changes
  - Additional TLV fields and associated management attributes / MIB amendments
  - No changes to the TLV processing state machines

New Features

Overview

- 5 New areas discussed
- Consensus to present progress on these items to TF for feedback
- All targeting TLV changes and only
- New areas
  - Watchdog – Lennart
    - Split out from Power Cycle
  - Unused Channel Power Loss Owner (Punch) – Fred
    - Additional baseline contribution in Fred's presentation
  - 10GBASE-T PoE – Yan
  - Dual PD Interfaces – Yan
  - Power Price – Bruce
    - Use of the term "Price" in this context cleared with 802.3 and 802.3bt Chairs
    - Additional background material included in Backup

Watchdog

- The power cycle feature can be extended with a timeout parameter. The intention is that the PSE would wait with the shutdown at least the time specified as timeout. A PD could send new frames, where the timeout would override the existing PSE timer. This would allow a watchdog like function. A PD should also be able to end this process.

Punch - Unused Channel Power Loss Owner

For the worst-case configuration, the PSE provides at least the minimum power level for a PD request—see table line one.

By default PSEs using LLDP own P_{unch}, and reduces the power provided so that the PD receives the power requested—see table line two.

Class 6 and 8 PDs may negotiate the ownership of the unused channel power. P_{unch} = 0 in the last table line.

Channel	Power Watts				
Length (m)	**PSE**	**Channel**	**PD**	**Comment**	**Punch (W)**
100	60.0	9.00	51.0	IEEE 802.3bt Type-3, Class 6, TBD	0.00
50.0	54.8	3.75	51.0	The PSE retains unused channel power	5.25
50.0	*60.0*	*4.50*	*55.5*	*The PD use all Pclass power, $P_{unch} = 0$*	0.00

Punch - Unused Channel Power Loss Owner

The existing power negotiation mechanism allocates P_{unch} indirectly. PDs request the power required.

Class 6 or 8 PDs may request up to P_{Class} from the PSE using LLDP. When $P_{PD} = P_{Class}$, $P_{unch} = 0$.

PSEs may provide worst-case power, P_{Class} from Table 33-7.

$P_{unch} = 0$ by default for Class 6 or 8 PDs when LLDP is not used.

The PD power value may be dynamically allocated.

Punch - Changes to PoE Clauses

No additional TLV changes

No MIB changes

No State Diagram change

10GBASE-T PoE – Feature Overview

- Problem Statement
  - Actual losses in the channel maybe less than worst case scenario described by the standard. Unused budget can be communicated to the MDI partner and recouped / reused
  - Measurements on PSE and PD may cost for extra designs
  - Lower cabling losses from better cabling
- Feature Overview (optional feature)
  - 10GBASE-T transmission requires Cat6 Cable which is better than Cat5 described in the spec.
  - If PSE and PD works on 10GBASE-T mode, PSE and PD can decide to use channel maximum DC pair loop resistance of Cat 6 for worst case channel losses.
- Feature Items for Discussion
  - Indicator that PSE and PD works on 10GBASE-T mode and/or better cabling
  - Coordination with work on “Unused Power Owner”

Dual PD Interface - Feature Overview

- Problem Statement
  - Certain applications may require two separate powering circuits inside a PD for two loads, such as a heater and a camera inside
    - Refer to option 2a in http://www.ieee802.org/3/bt/public/nov14/darshan_11_1114_rev_07.pdf in 2014 Nov. meeting.
- Feature Overview
  - PD has two separate loads behind a PD PI and can work simultaneously
  - Use L2 negotiation to recognize two separate loads behind one PD PI.
  - Use L2 to enable allocation of power on each pair-set
    - E.g. Communicate portion of overall power being requested for A vs. B
- Feature Items for Discussion
  - Indicate PD has two separate loads behind one PD PI;
  - Per pair-set power request/allocation.

Implementation examples for Option 2 PDs

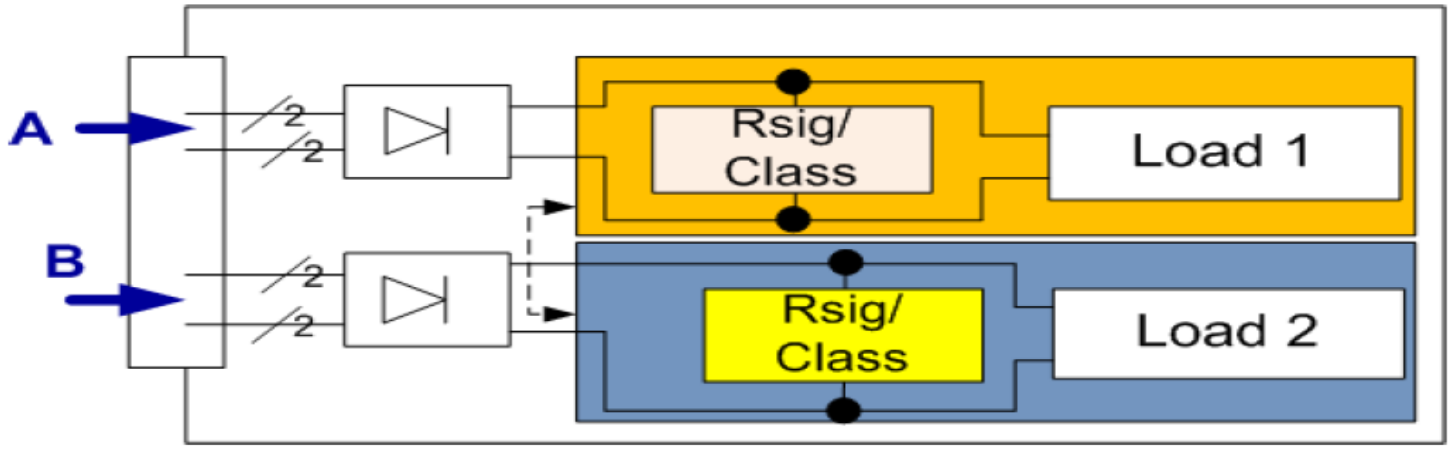

Details of option 2

Option 2a

-The PD was designed to work with 4P or 2P.
  -It has two loads.
  -**The <u>total power</u> of the two loads are Type 1 or Type 2 levels.**
  -Both loads can work simultaneously (No thermal issues. Ppd<Type 1/2)
  -The “Arrow” that indicates invalid detection signature on the un-powered pairs when connected, is disconnected now, to allow valid signature on Mode A and Mode B pairs so PSE knows that it is allowed for 4P operation.

Power Price - Proposal

- TLV **PowerPriceIndex** (16 bit unsigned integer)
  - Optional
  - Nominal value = 1000
  - Indicator of how available power is to the PSE
    - from external sources or internal battery
  - Higher value means power is less available
- Is set only by PSE
  - usually the same for all ports on a box; not required to be
- Used in conjunction with PDs requesting power allocation through existing mechanisms

Power Price - Proposal, notes

- Algorithms for how PSEs set PowerPriceIndex or how PDs use it do **NOT** need to be standardized and are out of scope for this standard
- End-use devices are shipped with default behaviors
  - Can be adjusted by installer or user as needed

- Example uses
  - PoE light begins to dim when PowerPriceIndex exceeds some threshold and turns off entirely when a second threshold is reached.
  - Total demand by PDs nears PSE power capacity so PSE increases PowerPriceIndex to stay under limit (and later decreases when capacity used drops)

Power Price - Power domains

- Links to sources (or sinks) of power
  - e.g. AC grid or local DC
- Other domains

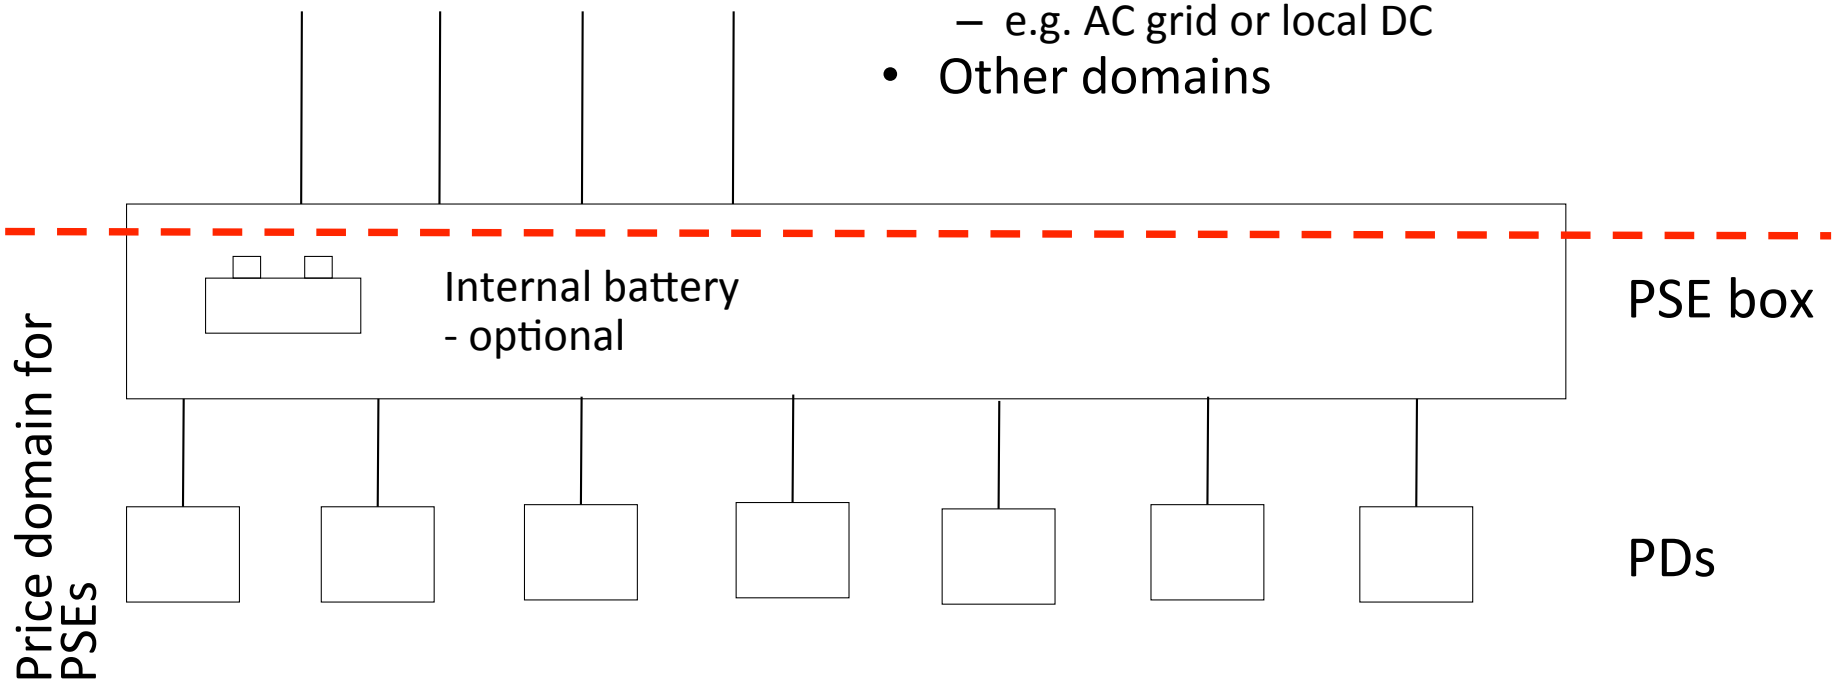

Next Steps

Next Steps

- Solicit feedback from TF on direction of the features

Backup

Power Price - Related Information

- Following slides provide context for how a price index can be used in future
- They do **not** affect the content of the standard

- For further information on Local Power Distribution
  - Bruce Nordman, bnordman@lbl.gov, http://nordman.lbl.gov

Power Price - Background

- **Traditional view**: Electricity always equally available across location and time
  - Utility tariffs were flat
  - Everything (almost) part of the same power domain
  - No (well, few) capacity constraints
  - So, no need to communicate about availability
- **Future**: Power availability varies – space and time
  - Dynamic utility rates
  - Local generation and storage; off-grid operation
  - **Local Power Distribution** – networked (within/between buildings); like local area networking

Batteries for building-scale rely on this for value

Power Price

- **Nanogrid**: Single domain of power; single voltage, capacity, reliability, administration, scarcity-index
  - PoE Switch **is** a nanogrid controller
- Nanogrid controllers may contain storage (battery)
  - Storage makes possible network models; comm. & power
- Nanogrid controllers will negotiate with other controllers and local generation to exchange power

- Power Price Index is how entities know value of power
  - How to produce power, manage storage, exchange, consume
  - Eventually will include (non-binding) forecast, but for now just current state
  - Local to each nanogrid

Power Price - Other Issues

- Prices are expressed in currencies and units of energy (usually kWh). Some other mechanism may be needed to convey what currency and price corresponds to nominal value (1000). This would be constant for a given link establishment.
- Ideally there would be a price and a (non-binding) forecast of future prices, with time expressed in seconds relative to ‘now’. Such a series of pairs of values may be more data than is reasonable to put into LLDP.

Power Price - Scaling structure: communications and power