LLDP / LLDP-MED Proposal for PoE Plus

John Jetzt (Avaya)
Peter Blatherwick (Mitel Networks)
Dan Dove (ProCurve Networking by HP)

Supporters and Contributors

Dan Romascanu; Avaya: co-author ANSI/TIA-1057 (LLDP-MED)

John Jetzt; Avaya

Meir Deutsch; Avaya

David Frattura; Enterasys: co-author ANSI/TIA-1057 (LLDP-MED)

Peter Blatherwick; Mitel: editor ANSI/TIA-1057 (LLDP-MED)

Dan Dove; ProCurve

Manfred Arndt; ProCurve: co-author ANSI/TIA-1057 (LLDP-MED)

Paul Congdon; ProCurve: editor 802.1AB-2005 (LLDP)

Scope

- Propose an efficient Layer-2 power management protocol for enhanced power allocation, that includes the following:
  - Fine-grain power negotiation
  - Ongoing dynamic re-negotiation (e.g. video call in process)
  - Power priority (e.g. must keep “red phone” alive)
  - Backup power conservation (e.g. extend UPS battery life during disasters)
- Must be practical for both cost-restrained and feature rich endpoints
- Must be backwards compatible with existing pre-PoE Plus devices

Objectives

- Widespread industry adoption
  - Simple and leveraged design increases chances of vendor adoption
  - Simple design leads to low development cost
- Interoperability with many endpoint device types
  - Low complexity ⇔ higher interoperability potential
- Must be compatible with existing pre-PoE Plus implementations
  - Basic operation without 802.3at is required
- Allowance for low-cost, highly constrained devices
  - KISS !!
  - Support “basic” static vs. extended “semi real time” dynamic negotiation
- Easily extensible for future needs

Endpoint Device Considerations

- PoE Endpoint device types -- some examples
  - IP Phones, video phones and other communication devices
  - WLAN Access Points, video cameras, Pan-Tilt-Zoom video cameras
  - Card access readers (e.g. door entry)
- Fundamentally, there are two classes of devices:
  - Steady state power usage (WLAN APs, video cameras, card readers)
  - Dynamic usage (IP Phones, PTZ video cameras)
- Steady state devices require little additional functionality, so features like standby/dynamic usage should be optional for PDs with reduced complexity
  - Max power required / supplied is sufficient for most “steady state” devices
- Some devices, like PTZ security cameras, may change power usage very frequently and also require near instantaneous responsiveness
  - Dynamic negotiation for rapidly fluctuating devices seems impractical and likely will negatively impact performance and interoperability (guarantee max allocation)

Consideration for low cost / highly constrained endpoints is CRITICAL !

LLDP Overview

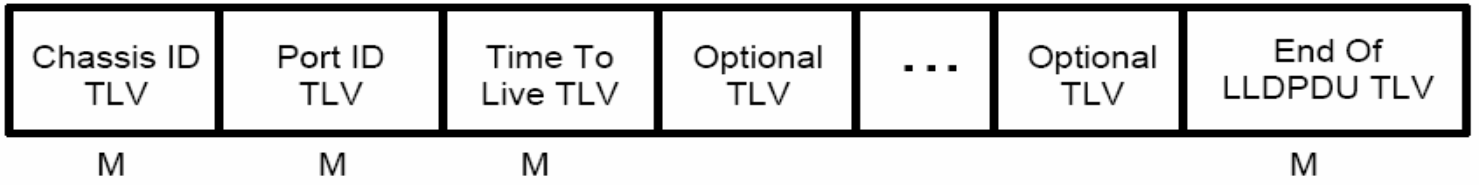

- Basic Functions (IEEE 802.1AB-2005*):
  - Simple one-way neighbor discovery protocol with periodic transmissions
  - LLDP frames are not forwarded, but constrained to a single point to point link
  - LLDP frames contain formatted TLVs (type, length, value)
    - Globally unique system and port identification
    - Time-to-Live information for ageing purposes
    - Optional system capabilities (e.g. router, IP phone, wireless AP)
    - Optional system name, description, and management address
    - Organizational extensions
  - Receiver stores information in a neighbor database, accessible via SNMP MIB
  - Receiver ages MIB to insure only valid network data is available
  - Management applications can harness the power via SNMP

* Currently planned for revision, already approved to submit PAR in Nov ’06

LLDP TLV Extensibility

Easy to define organizational extensions

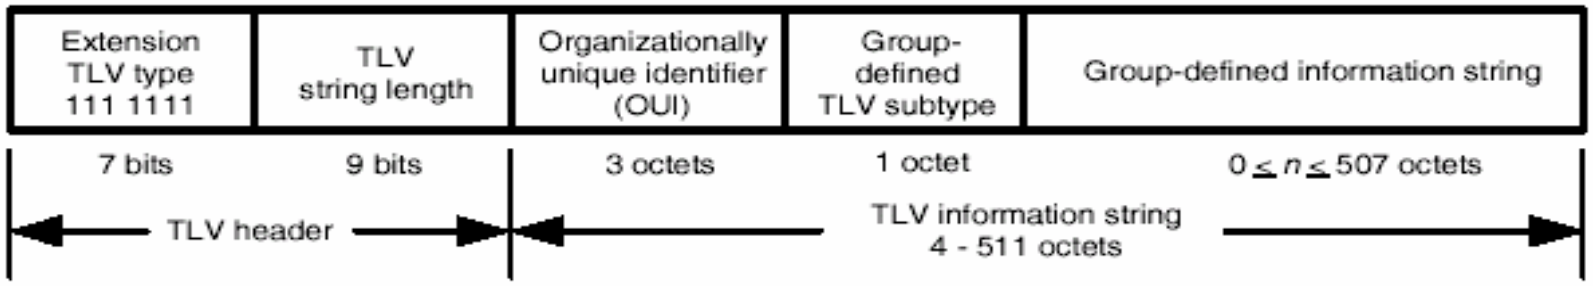

Figure 9-12—Basic format for organizationally-defined extension TLVs

There are currently three organizational extensions:

1. IEEE 802.1
   - Port VLAN, Port & Protocol VLANs, VLAN Name, Protocol Entity
2. IEEE 802.3
   - MAC/PHY configuration, Power, Link Aggregation, Maximum Frame Size
3. TIA, LLDP-MED

LLDP-MED Advanced Power Management

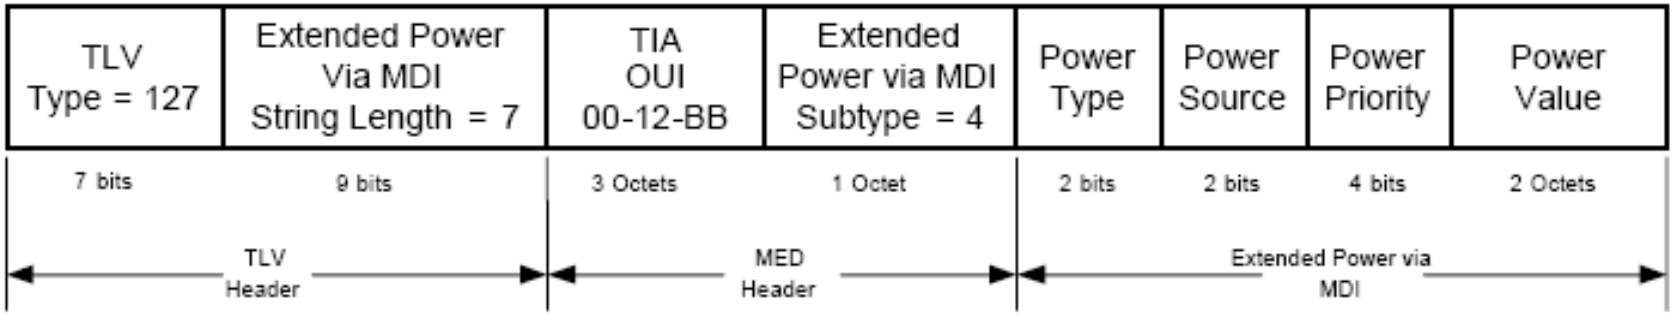

Figure 12 – Extended Power-via-MDI TLV Format

- Allows PoE devices to advertise their power related information
  - Setting of power management policy is outside the scope of LLDP-MED
  - PoE usage policy based on this data; however, would be appropriate for 802.3at
- Values
  - Power Type: PSE or PD
  - Power Source: Normal or back-up power conservation
  - Power Priority: Critical, High, Low
  - Power Value: 0 – 102.3 Watts (0.1 W increments)

Is LLDP a Stateless Protocol?

IEEE 802.1AB-2005:

- LLDP updates are limited to no faster than 1 per second
- Intended as one-way advertisements, without req / acks
  - Dramatically simplifies implementations
  - Bounds performance requirements for scalability
- On occasion, the management entity might see a change in state of a peer and perform some local database maintenance operation
- By definition, anytime a local value changes, LLDP sends a frame thus triggering a packet in reverse (limited to once per second)
- Packet loss between PD and PSE is negligible, recovered on next advertisement
- The protocol itself is stateless, but the management entity above is required to maintain state and may act on information from peers

Start-up / State Interaction (example)

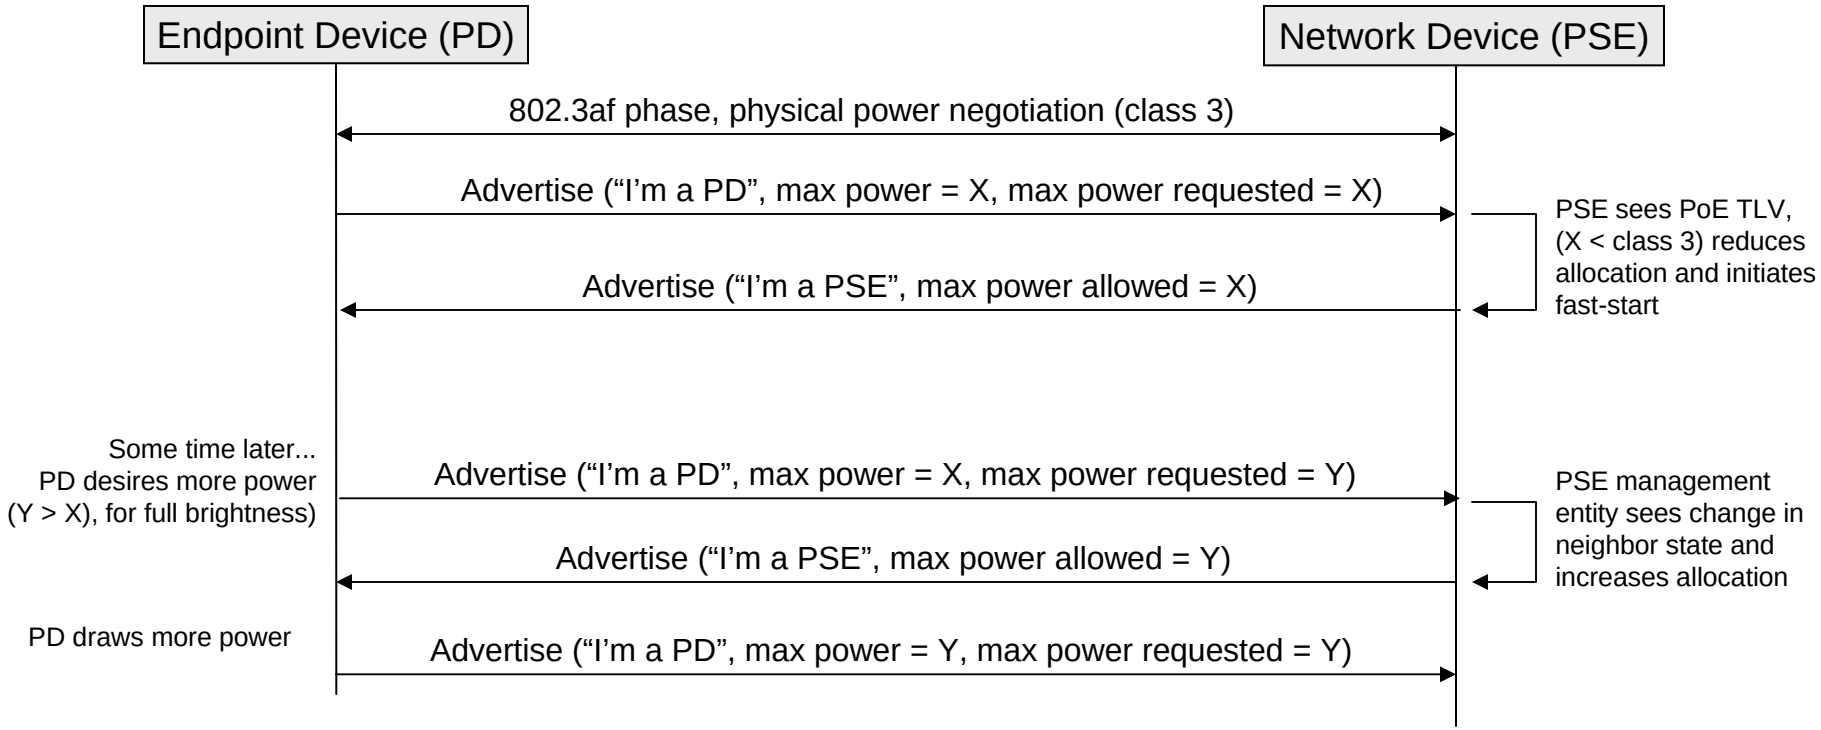

Usage Rules:

- PD shall never request more power than physical 802.3af class
- PD shall never draw more than max power advertised by PSE
- PSE may deny any PD drawing more power than max allowed by PSE
- PSE shall not reduce power allocated to PD, that is in use
- PSE may *request* reduced power, via conservation mode

TLV Values to Consider

- Fine grained max power required / available
  - Well above current 802.3at targets
- Power conservation modes
- Power priority (critical, high, low)
- Optional dynamic power negotiation (extended devices)
- Support both initial booting and ongoing operation
- Actual average power ?!
  - Seems more appropriate for PSE to monitor this
  - Many PSE chipsets report actual power today
  - Endpoint should not be required to maintain that level of accuracy anyway

Advantages of LLDP-MED

- Existing and well defined standard
  - Rapid development, deployment and easily understood
  - Supported by existing implementations, tools and processes, already deployed
- Extended Power-Via-MDI TLV covers most, if not all, requirements today
  - Fine grained power reporting, 0.1 increments, up to 102.3 W max
  - Power source (e.g. switch running on backup), allows for PD to go to power conservation mode
  - Power priority, to influence PSE policy and detect mismatches (e.g. must keep "red phone" alive)
  - New requirements (if any) are incremental, and readily extended
- High interoperability potential and industry acceptance
  - Simple / decoupled / stateless design ⇔ high chances of working correctly
  - Based on pre-existing design and code base, already debugged
- Easy, low cost implementation
  - Simplistic design is easy to work with and readily applicable to low cost / constrained devices
  - Code base already shipping in many important cases (notably IP Phones)!
  - No additional protocol required (eliminates redundant effort and significant complexity)

Possible LLDP Issues

- Possible extensions:
  - Requested maximum power (over and above currently required), PD
  - Available maximum power (over and above currently supplied), PSE
  - Can be written either as part of 802.3at standard or 802.1AB maintenance

- Timeliness of state updates
  - Is the once a second throttling of 802.1AB an issue (questionable) ?!
  - Motion to create a PAR for an 802.1AB revision has been approved.
    - A "fast start" capability has been discussed as a possible item for the revision
  - Fast start may improve timeliness of neighbor state changes
    - Fast start has benefits beyond 802.3at's needs (e.g. AVB and CM discovery)

Summary

- LLDP-MED provides several advantages for PoE Plus
  - Existing, well defined standard
  - Covers most, if not all pragmatic needs
  - Simple and effective
  - Industry accepted solution, already deployed
  - Reduced complexity and implementation cost, critical for cost-restrained endpoints
  - Easily extensible for future needs
- Stateful protocol does not seem to be required
  - Timeliness needs can be met without statefulness at protocol level
  - Packet loss is not a significant issue, does not justify added complexity

Recommend LLDP-MED as starting point

Next Steps

- Incorporate new 802.3at TLV extensions and a fast start capability into the 802.1AB revision PAR, if all parties are in agreement to proceed
- Clarify PoE timeliness requirements as input to LLDP fast start update
- Define LLDP-MED extensions to meet technical requirements (if any)
- Define behavior rules within 802.3at, that describe the management entity interactions when PoE TLVs are sent and received