

# Extending Link Layer Discovery Protocol for Larger Databases

Paul Bottorff

Aruba a Hewlett Packard Enterprise Company

[Paul.Bottorff@hpe.com](mailto:Paul.Bottorff@hpe.com)

Paul Congdon

Tallac Networks, Inc.

[paul.congdon@tallac.com](mailto:paul.congdon@tallac.com)

## Abstract

*The Link Layer Discovery Protocol (LLDP, IEEE Std 802.1AB [1]) is widely used throughout the network industry in enterprise, data center, and carrier networks for discovering links, stations, port characteristics, network topologies, and device attributes. LLDP achieves discovery by exchanging copies of a database with attributes represented by type-length-value (TLV) records. However, the current LLDP database size limitations prevents LLDP from supporting more demanding network discovery applications such as router discovery. Currently, LLDP limits the size of the information advertised from a port to a single frame of information. Here we propose a protocol for extending LLDP to support multi-frame databases allowing a port to advertise over 100 Kilobytes of information. The technique proposed for extending LLDP is backward compatible with existing LLDP implementations to allow easy deployment in existing networks.*

## Introduction

The Link Layer Discovery Protocol (LLDP) operates over network links to build a distributed database with information from all the neighbor ports attached to each link. Each Link Layer Discovery Protocol Data Unit (LLDPDU) used to advertise information is encoded within a single frame as a series of Type Length Values (TLVs). The TLVs exchanged by LLDP contain information advertised from the originating port. The LLDP TLV's formats and usage are defined by multiple organizations which fall into three categories: 1) the IEEE 802.1 committee, 2) other standards organizations, 3) any other organization or product vendor using an IEEE 802 Organizationally Unique Identifier (OUI).

Currently the IEEE 802.1 and 802.3 committees define 30 TLVs which LLDP can advertise from a port. These include the basic set defined in the LLDP standard IEEE Std 802.1AB-2016, additional TLVs defined in IEEE Std 802.1Q-2018, and TLVs defined in IEEE Std 802.3-2018. Some of these TLVs are variable in size and some can have multiple copies in the databases. If we size a database containing a single copy of the maximum sized TLV for all current 802.1 and 802.3 TLVs, the size of the LLDPDU would be 2533 octets which is larger than fits in the maximum LLDPDU for a common Ethernet frame (i.e. 1500 bytes for a basic Ethernet without VLAN tagging). These numbers do not consider TLVs with multiple copies, TLVs defined by other standards organizations or TLVs defined by other organizations. Current LLDP implementations work around these limitations by not including all the TLVs and rarely use maximum sized TLVs, however, applications for LLDP continue to grow. The information exchanged will eventually require LLDP extensions or force users to support multiple incompatible discovery protocols.

A major application for an extended discovery protocol is discovering router-to-router interfaces along with their network addressing and characteristics. This application is driven by the rise of new routing protocols currently under development at the Internet Engineering Task Force (IETF) such as Link State Vector Routing (lsrv). Though traditional link state routing protocols such as Open Shortest Path First (OSPF) and Intermediate System to Intermediate System (IS-IS) have built in discovery protocols, other routing protocols such as Boarder Gateway

Protocol (BGP) do not. Given the broad deployment of LLDP an extended LLDP with sufficient database capacity would allow LLDP to support discovery for these new routing protocols and enable consolidation to a single discovery protocol.

Another application for an extended LLDP protocol is support for Time Sensitive Networks (TSN) where the network imposes a small maximum frame size to limit latency for time sensitive traffic. For these TSN networks the amount of information LLDP can exchange can be highly constrained since current LLDP can only advertise a single frame of information from a port even if the single frame is restricted to a small size, such as 64 bytes. The ability to advertise LLDP information using multiple frames would allow these LLDP applications to use LLDP without highly constrained LLDP database sizes.

## **New Discovery Protocols**

Currently there are multiple new discovery and configuration protocols at the IEEE and IETF. These are targeted at specific applications for data center virtualization, stream registration, and router discovery.

To support virtualization in data centers the IEEE has standardized the Virtual Station Interface Discovery and Configuration Protocol (VDP, Clause 41 of IEEE Std 802.1Q-2018, IEEE Std 802.1Qcy-2019 [4]) which is used to exchange information about a virtual interface within a server with an edge bridge or router. The VDP protocol is specifically targeted at data center server virtualization and allows discovery, configuration, state, and liveness of virtualized interfaces of virtual machines residing within a server. The VDP protocol may scale to any number of Virtual Station Interfaces (VSIs).

To support the requirements of industrial automation time sensitive networks, the IEEE is standardizing the Link-local Registration Protocol (LRP, IEEE P802.1CS [5]) and the Resource Allocation Protocol (RAP, IEEE P802.1Qdd [6]). The Link-local Registration Protocol provides the ability to replicate a database optimized for 1 Megabyte of data over a data link. The Resource Allocation Protocol runs as a distributed application over the Link-local Registration Protocol to support large stream reservation databases for industrial applications.

At the IETF work on new link state protocols for data centers such as Link State Vector Routing (lsvr) has resulted in new discovery protocol proposals: the Layer 3 Discovery and Liveness protocol [2]; the BGP Neighbor Discovery protocol [3].

These new discovery protocols from the IEEE and IETF are very capable at their target applications, however are more complex than required for simple discovery applications. LLDP on the other hand is a general purpose discovery protocol which can be used for a wide variety of applications and is already very broadly deployed.

## **Objectives for an Extended LLDP Protocol**

The IEEE has approved a project to extended LLDP (P802.1ABdh [7]) which intends to:

- Extend LLDP to support transmission and reception of a set of Link Layer Discovery Protocol (LLDP) Type Length Values (TLVs) that exceed the space available in a single frame.
- Maintain existing functionality while communicating with a peer LLDP agent that supports updated functionality.
- Provide a reasonable sized LLDP database for time sensitive networks with restricted LLDP Data Unit (LLDPDU) size.

For the protocol proposed in this paper we consider some clarifications and additional objectives which are motivated by ongoing work in other committees and by general protocol considerations. These include:

- Support advertisement of multi-frame databases from each link neighbor of over 100K bytes when using maximum Ethernet sized frames.
- Support exchange of a single-frame database from each link neighbor with any previous LLDP version.
- Support multi-frame databases of small frames where the chassis+port+ttl+extension TLVs fit in a single frame and where the size of the advertised database can be constrained by the frame size.
- Support both shared and pt-pt media.
- Operate on links where MAC addresses are duplicated.
- Minimize multicast use, periodic transmissions, and database updates.
- Support pacing of data by receivers.

The P802.1ABdh objectives leave open the question of how big the supported database needs to be. Looking at the requirements for the router discovery application in the IETF, a maximum database size of around 64K octets is required. The 64K requirement could exist for each protocol stack supported, so a dual stack application could support up to 128K octets of TLVs.

The P802.1ABdh project objectives require backward compatibility with existing LLDP implementations. The objective is that a new P802.1ABdh implementation is able to exchange a single frame of information with an existing IEEE Std 802.1AB implementation without any changes in the existing LLDP implementation. The existing LLDP implementation will silently ignore any of the additional TLVs provided in the extended database. As a consequence, an existing LLDP implementation will not be able to exchange the extended database.

The requirement for supporting small maximum frame sizes has many unexplored consequences. First, it is important to understand how large the required database will be for this application since small frame sizes will constrain the maximum database size even with multiple frame support. There will likely be a limit to the number of frames that can be used to transfer the database, and if those frames are constrained in size, the overall database will be constrained in size. In addition, we need to understand if there are any constraints on the size of TLVs for this use case. Current LLDP TLVs have a maximum information field length of 511 octets and if this TLV size is too large to meet the timing constraints, then alternative solutions need to be considered (e.g. TLV fragmentation). Without further understand of these requirements, this proposal considers the simplest and most compatible approach to these constraints. The constraints which result include:

- Individual TLVs may not span multiple frames and can not be larger than the minimum frame size
- The chassis+port+ttl+extension TLVs must fit in a single frame
- The number of additional frames allowed in the database will be constrained along with frame size

The current LLDP protocol is a multicast protocol, however for support of larger databases it is desirable to use a unicast protocol along with receiver data pacing to prevent overflow. This will avoid the historic problems seen in protocols such as IS-IS and OSPF which use multicast to build distributed databases.

Another historic problem with multicast protocols is the required periodic transmission results in background traffic which can overloading station control processors. To prevent these overloads it is desirable to only transmit/update when there are changes in the database. Also it is desirable to only transmit the frames which require updating.

For router discovery applications two additional requirements are being considered by the IETF. These are authentication of the database and liveness. The LLDP protocol makes periodic updates and supports a time to live value for the information, however this is at a very slow rate and results in a poor liveness indication. The LLDP protocol is not well suited to frequent keep alive exchanges, however alternative protocols exist. Rather

than use LLDP for liveness the IEEE Connectivity Fault Management protocol (CFM, IEEE Std 802.1Q-2018 cl 18-22 [8]) already has the ability to implement very capable keep alives.

Authentication for LLDP is a topic which needs focused study by security experts. As such it is more rational for IEEE to launch a separate security project where the topic can be handled in depth rather than try to add authentication to the exiting P802.1ABdh project.

## Current LLDP Operation

LLDP is a layer 2 protocol that allows a station to advertise capabilities and state in the form of a local database at each MAC attached to a link. Each LLDP agent uses the information advertised by other LLDP agents to build a database composed of remote information advertised by the other agents along with its own local information.

An LLDP (IEEE Std 802.1AB-2016 [1]) agent transmits information associated with a local MAC periodically (figure 1, A). The transmission is typically to a multicast address (unicast is supported for Wi-Fi networks) which is selected from those multicasts that do not pass through 802 bridges. Transmission begins when the local MAC is enabled and continues periodically (figure 1, A and B). Rapid updates are performed whenever some information in the local database changes (figure 1, C). The receiving LLDP agent compares the contents of the received TLVs with the previous version, and if anything differs, the entire previous database is discarded and replaced with the database represented by the newly received TLVs.




Figure 1: IEEE Std 802.1AB-2016 Link Layer Discovery Protocol Overview

## Extending LLDP

The basic idea used to extend the 802.1AB-2016 protocol to multiple PDU databases is to encode a description of the extension PDUs within the initial frame which will be used by the extended LLDP to retrieve the rest of the database. We call the initial frame, which is the single frame exchanged by the 802.1AB-2016 protocol, the LLDP Foundation PDU or F-PDU. The F-PDU is an 802.1AB-2016 LLDPDU which includes an additional TLV, called the Manifest TLV, that describes the extension PDUs (X-PDUs).

The Manifest TLV contains a unique identifier for each X-PDU that changes whenever any change is made to an X-PDU. Since changes in the extended database results in a change of the Manifest TLV they also result in a change in the F-PDU. The change in the Manifest TLV results in a rapid update of the F-PDU and starts the extended LLDP

(XLLDP). XLLDP uses the Manifest TLV to form requests for transmission of the needed X-PDUs.

If an 802.1AB-2016 implementation receives an F-PDU with a Manifest TLV it simply ignores the Manifest TLV since it is an unrecognized TLV type. An 802.1AB-2016 implementation therefore operates without the X-PDUs, however maintains its existing single LLDPDU operation.

XLLDP is a unicast request response protocol where the XLLDP receiver requests each changed or missing X-PDU. XLLDP only makes requests for X-PDUs that are required to update its database, therefore X-PDU requests are only made to build the initial database or when an X-PDU is changed or added to the database. XLLDP uses the X-PDU descriptions listed in the Manifest TLV to determine if it needs to update an X-PDU by comparing the Manifest TLV list with the current database.

The unicast addresses used by XLLDP to make requests and responses are those used as the source MAC address in the frame which encodes the F-PDU. Additional addressing is provided on each request by specifying the destination chassisID+portID rather than the source chassisID+portID used in a conventional LLDPDU. Response X-PDUs encode the source chassisID+portID as a normal LLDPDU transmission.

Each extension request PDU specifies a list of X-PDUs for the source to transmit. The requester can use the requests to pace transmission by controlling the number and spacing between requests and can request re-transmission if desired. The XLLDP allows only a single outstanding request to each remote LLDP agent.

## New PDU Formats

Three new types of PDUs are used by the Extended Link Layer Discovery Protocol (XLLDP). These are the Foundation PDU (F-PDU), the Extension PDU (X-PDU) and the Extension Request PDU (XREQ-PDU).

### Foundation PDU (F-PDU)

Figure 2 shows the format for the F-PDU. The F-PDU is an LLDPDU as specified by 802.1AB-2016 and is identified by the LLDP Ethertype as specified in 802.1AB-2016. An F-PDU is a particular type of LLDPDU which contains a Manifest TLV describing a list of X-PDUs. The Manifest TLV immediately follows the TTL TLV within the F-PDU. An 802.1AB-2016 implementation which receives an F-PDU will simply ignore the Manifest TLV and process the remainder of TLVs as normal.




Figure 2: Foundation PDU (F-PDU) Format

The Manifest TLV format shown in Figure 3 is identified by a new base LLDP TLV type. The Manifest TLV contains a 7 bit field with the number of X-PDUs described by the Manifest TLV followed by an array of n 6 octet Extension PDU descriptors. The number of X-PDUs is indicated by a separate field from the TLV length to allow implementations to use a fixed size Manifest TLV with a variable number of Extension PDU descriptors. If the number of Extension PDUs is less than the TLV string length, then the first n Extension PDU descriptors are valid.

Each X-PDU descriptor in the Manifest TLV provides a unique identifier for the X-PDU. The identifier is composed of 3 fields which are a 7 bit X-PDU identifier, an 8 bit X-PDU revision, and a 32 bit X-PDU check value. The X-PDU number along with the X-PDU revision provide a 16 bit identifier which uniquely identifies most X-PDUs.

Encoding the X-PDU number and revision within the Manifest TLV allows insertion and deletion of X-PDUs

without updating the other X-PDUs. The check value is used to cover boundary cases where the X-PDU number and revision fail to uniquely identify the X-PDU. The check value bits are calculated using all octets of the X-PDU in an MD5 hash calculation and then taking the low order 32 bits of the result. An example where the check value would be used is the case where an X-PDU number is re-used and the X-PDU revision is restarted rather than updated, then the X-PDU check along with the new random X-PDU revision provide uniqueness for the new entry within an error of about 1 in a trillion. It is expected this condition would only happen when the XLLDP agent loses its state for instance after a power-on reset.




Figure 3: Manifest TLV Format

### Extension PDU (X-PDU)

Figure 4 shows the format of an Extension PDU (X-PDU). An X-PDU is identified by an XLLDP Ethertype plus an XLLDP subtype. This Ethertype is different from the LLDP Ethertype to assure the XLLDPDUs are not mistakenly used by an 802.1AB-2016 implementation. The subtype is used to identify the XLLDPDU as an X-PDU. The subtype is included to allow for future PDU types and the potential use of the Ethertype by other protocols with similar characteristics. There is also a XLLDP revision number which is intended for future revisions of XLLDP.

Each X-PDU contains a ChassisID TLV and PortID TLV which identify the source for the X-PDU. There is no time to live TLV in an X-PDU since the TTL is set from the F-PDU which established the database. Following the ChassisID and PortID TLVs is an Extension PDU Identifier TLV (XID TLV) used to identify this X-PDU. The Extension Identifier TLV can be followed by any number of optional TLVs that fit within the frame.




Figure 4: Extension PDU (X-PDU) Format

The X-PDU Identifier TLV (Figure 5) uniquely identifies this X-PDU within the database addressed by the ChassisID+PortID using an X-PDU number and an X-PDU revision. The 7 bit X-PDU number identifies this X-PDU from the other X-PDUs local to this XLLDP agent. The 8 bit X-PDU revision indicates the revision since database initialization. The X-PDU revision starts at a randomly selected number on reset and is incremented modulo 256 whenever any change is made to the PDU. At reset the X-PDU revision number starts at a new randomly selected number. The 32 bit check value, used in the Manifest TLV to verify the unique identification of the X-PDU, is not carried in the X-PDU and therefore is calculated when a new X-PDU is received by the XLLDP agent.




Figure 5: Extension PDU Identifier TLV (XID TLV) Format

### Extension Request PDU (XREQ-PDU)

Figure 6 shows the format of the Extension Request PDU (XREQ-PDU). An XREQ-PDU is identified by the XLLDP Ethertype plus an XLLDP subtype. This Ethertype is different from the LLDP Ethertype to assure the XREQ-PDUs will not be used mistakenly by an 802.1AB-2016 implementation. The subtype is used to identify the XLLDP PDU as an XREQ-PDU. There is also a XLLDP revision number which is intended for future protocol revisions.

Each XREQ-PDU contains a ChassisID TLV and PortID TLV which identify the destination for the XREQ-PDU rather than the source like a normal F-PDU or X-PDU. There is no time to live TLV in an XREQ-PDU. Immediately following the destination ChassisID and PortID TLVs is an XREQ-TLV which provides a list of the wanted X-PDUs. No other TLVs except the (optional) End of LLDPDU TLV are present in the XREQ-PDU.




Figure 6: Extension Request PDU Format

Each XREQ-PDU contains an XREQ-TLV (Figure 7). The XREQ-TLV contains a list of X-PDU descriptors which the sender of the XREQ-PDU needs updated. Each descriptor element provides the X-PDU number, revision and check value for the needed X-PDU. The number of X-PDUs descriptors encoded in the XREQ-TLV is determined by the length of the TLV. Each XREQ-TLV also includes a 16 bit request number which the sender increments modulo 64K for each new request issued. On reset the request number begins at a randomly selected starting point.




Figure 7: Extension PDUs Request TLV (XREQ TLV) Format

## Extended Link Layer Discovery Protocol (XLLDP)

Figure 8 illustrates XLLDP operation. Any LLDP exchange including an XLLDP exchange begins with the LLDP LLDPv2 Proposal

protocol as specified in IEEE Std 802.1AB-2016. LLDP is used to establish all databases and to exchange the F-PDUs which describe the X-PDUs extending each database to multiple PDUs (Figure 8, G). XLLDP never creates new databases, however is used to transfer the X-PDUs which extend the database established by LLDP.

XLLDP does not run periodically, but rather is triggered by the receipt of an F-PDU containing a Manifest TLV with descriptors for X-PDUs. Missing or revised X-PDUs are identified by comparing the Manifest TLV X-PDU descriptions with the current database. XLLDP is responsible for obtaining any required X-PDUs and updating the its copy of the remote database. Once the database is up to date XLLDP will not run again until it determines through a received Manifest TLV that the database needs to update some X-PDUs. Since the Manifest TLV is transmitted in the F-PDU by LLDP, any change to an X-PDU will be result in a change to the Manifest TLV and therefore a change in the F-PDU (Figure 8. A). LLDP will therefore immediately update the F-PDU whenever any X-PDU is changed in turn resulting in XLLDP running to request the required X-PDUs (Figure 8, B).

Once the required X-PDUs are determined by inspection of the changes to the Manifest TLV within the F-PDU, XLLDP will transfer the list of required X-PDUs identified from the Manifest TLV. XLLDP executes until it has retrieved all the needed X-PDUs or determined they are no longer required.

To transfer the X-PDUs XLLDP makes explicit requests for each needed X-PDU to the source XLLDP agent (Figure 8, C). These requests are encoded in XLLDP XREQ-PDUs (Figures 6, 7). Each XREQ-PDU can request multiple X-PDUs. Multiple XREQ-PDUs can be used by XLLDP to transfer the needed X-PDUs (Figure 8, C, E). XLLDP only allows a single pending XREQ-PDU for each peer XLLDP agent at a time. Since it is possible XLLDP has an outstanding XREQ-PDU to multiple peers XLLDP agents it may further limit the transmission of XREQ-PDUs to pace its receiver by limiting the number of peer agents with pending XREQ-PDUs. XLLDP controls pacing of frames by the number of outstanding X-PDU requests and the time interval between these requests. XLLDP times out each XREQ-PDU. If no response is received after the timeout XLLDP can retry the request. After retrying, XLLDP will assume the X-PDU is unavailable and will invalidate the X-PDU in its database, log an error, and continue to recover all the X-PDUs that are available at this time.

Each XREQ-PDU is transmitted in a unicast MAC frame addressed to the MAC address provided as the source address of the frame carrying the F-PDU. The XREQ TLV carries a list of the descriptors from the Manifest TLV for the requested X-PDUs.

The receiver of an XREQ-PDU uses the ChassisID and PortID of the XREQ-PDU to determine what local database contains the requested X-PDUs (Figure 8, C&D). If the receiver does not have the database matching the ChassisID and PortID of the XREQ-PDU it discards the XREQ-PDU. If the receiver has the database it locates the X-PDUs that match the XREQ TLV descriptors and transmits the X-PDUs (Figure 2) to the source MAC or the XREQ-PDU making the request. If the receiver has the database, however does not have the matching X-PDUs it initiates an update of the F-PDU.

To suppress duplicate responses the sender keeps track of the last request number from the most recent XREQ TLV. If a new XREQ TLV with a request number matching the last response is received the sender can discard the XREQ PDU and not send any of the requested X-PDUs.




Figure 8: Extended LLDP (XLLDP) Operation

## XLLDP Operation Over Links with Duplicate MAC Addresses

In most cases 802 links don't have duplicate MAC addresses, however duplicates are possible under some circumstances. LLDP is designed to operate correctly even when duplicate MAC addresses are present on the link, and like LLDP, XLLDP is also resilient when operating on links with duplicate MACs.

To best describe the operation of XLLDP when duplicate addresses are present on the link, consider a classic CSMA/CD Ethernet example (Figure 9) with three cases of duplicate addresses:

- The switch sourcing an X-PDU can have multiple MACs connected to the link each with the same MAC address (Figure 9, switch A);
- The switch requesting an X-PDU can have multiple MACs connected to the link, each with the same MAC address (Figure 9, switch B);
- Some switch, which is not a party to an X-PDU exchange, may have MACs connected to the link with duplicates of the MAC addresses of switch A or switch B (Figure 9, switch C).




Figure 9: Example CSMA/CD Ethernet with Duplicate MAC Addresses

Figure 9 illustrates the example of a shared media Ethernet link connecting three switches. The switches each have 2 MACs attached to the Ethernet. The switch A uses the same MAC address "a" on both ports 1 and 2 connected to the Ethernet link.

The switch B uses the same MAC address “b” on both of its ports 1 and 2 connected to the Ethernet link. The switch C has two ports 1 and 2 connected to the Ethernet with MAC address duplicates of the Switch A and B addresses. Any frame transmitted on this Ethernet link is delivered to all MACs attached to the Ethernet. Within a switch incoming frames are steered to the Higher Layer Entities, including LLDP and XLLDP, as described in IEEE Std 802.1Q-2018 subclause 8.5.3. In our example frames addressed to “a” and “b” will be directed to the Higher Layer Entities of switches A, B and C where they will be further decoded based on the LLDP and XLLDP Ethertypes. An XLLDP transmission from switch B to destination address “a” will be delivered to the XLLDP agent on switch A ports 1 and 2 and to the XLLDP agent on switch C port 1. An XLLDP transmission from switch A to destination address “b” will be delivered to the XLLDP agent on switch B ports 1 and 2 and to the XLLDP agent on switch C on port 2.

For example an XLLDP agent operating in switch B port 1 of the figure 9 network sends an XREQ-PDU for an X-PDU of the switch A port 1 database. Switch B sends the XREQ-PDU in a frame to MAC destination address “a” since this was the source address used by switch A LLDP to deliver the F-PDU which caused the request. The XREQ-PDU sent from switch B will then be received by the XLLDP agents for switch A ports 1 and 2 and switch C port 1. On switch C the XLLDP agent will look at the destination ChassisID and PortID contained in the XREQ-PDU and determine these are not for any of its ports and therefore will discard the XREQ-PDU. On switch A (figure 10) a copy of the XREQ-PDU will be delivered to the XLLDP agent for both ports 1 and 2. Since the XLLDP agent will begin by checking the destination ChassisID and PortID the XLLDP agents will determine the XREQ-PDU is intended for the port 1 database. The XLLDP agent at port 1 checks the XREQ TLV request number against the last response it sent. If a response has already been sent for this XREQ XLLDP deletes the request. Otherwise, the XLLDP agent transmits the requested X-PDUs and records the XREQ TLV request number for future reference.

In our example the XLLDP agent operating in switch A port 1 will send the X-PDUs requested by the XREQ-PDU. These will be sent to the source address of the XREQ-PDU which in our example is MAC address “b” and therefore the X-PDUs will be received at the XLLDP agent for switch B ports 1 and 2 and switch C port 2. The XLLDP agent for switch B looks at the source ChassisID and PortID in the X-PDU and updates the remote databases which match the X-PDU ChassisID and PortID (figure 11). The X-PDU is also delivered to switch C on port 2. If switch C already has a remote database established by LLDP with the ChassisID and Port ID of the X-PDU, then the XLLDP agent updates that remote database.




Figure 10: XLLDP Example Switch A




Figure 11: XLLDP Example Switch B

## Summary

The LLDP protocol can be extended to allow advertisement of local databases of over 100 Kbytes while retaining backward compatibility with existing single frame implementation. The extension of LLDP allows implementations to support more demanding discovery applications such as the discovery of router neighbors and their addressing.

The proposed extension method uses an Extension Link Layer Discovery Protocol (XLLDP) to exchange Extended LLDPDU<sub>s</sub> (X-PDU<sub>s</sub>). XLLDP is based on a unicast rather than multicast protocol which enables pacing the transmissions of the databases by the receivers and elimination of periodic transmissions of X-PDU<sub>s</sub>.

## References

- [1] IEEE, IEEE Std 802.1AB-2016, Station and Media Access Control Connectivity Discovery, IEEE, 2016.
- [2] R. A. K. P. Randy Bush, "Layer 3 Discovery and Liveness," 02 11 2019. [Online]. Available: <https://datatracker.ietf.org/doc/draft-ietf-lsvr-l3dl/>. [Accessed 16 03 2020].
- [3] K. T. K. B. J. T. N. T. Xiaohu Xu, "BGP Neighbor Discovery," 26 11 2019. [Online]. Available: <https://datatracker.ietf.org/doc/draft-xu-idr-neighbor-autodiscovery/>. [Accessed 16 03 2020].
- [4] IEEE, IEEE Std 802.1Qcy-2019, Bridges and Bridged Networks - Amendment: Virtual Station Interface (VSI) Discovery and Configuration Protocol (VDP) Extension to Support Network Virtualization Overlays Over Layer 3 (NVO3), IEEE, 2019.
- [5] IEEE, "P802.1CS – Link-local Registration Protocol," IEEE, 17 02 2017. [Online]. Available: <https://1.ieee802.org/tsn/802-1cs/>. [Accessed 16 03 2020].
- [6] IEEE, "P802.1Qdd – Resource Allocation Protocol," IEEE, 27 09 2018. [Online]. Available: <https://1.ieee802.org/tsn/802-1qdd/>. [Accessed 16 03 2020].
- [7] IEEE, "P802.1ABdh – Support for Multiframe Protocol Data Units," IEEE, 05 09 2019. [Online]. Available: <https://1.ieee802.org/tsn/802-1abdh/>. [Accessed 16 03 2020].
- [8] IEEE, IEEE Std 802.1Q-2018 - IEEE Standard for Local and Metropolitan Area Networks—Bridges and Bridged Networks, IEEE, 2018.