

Topology Discovery in PROFINET

Iwan Schafer, Max Felser
Berne University of Applied Sciences, Engineering and Information Technology
Jlcoweg 1, CH-3400 Burgdorf
iwan.schafer@bfh.ch, max.felser@bfh.ch

Abstract

PROFINET, an automation network based on Ethernet, includes the automatic discovery of the physical topology with the Link Layer Discovery Protocol (LLDP) in its newest specifications. In this paper we present the objectives of LLDP in PROFINET and general information about the protocol and other parts of IEEE 802.1AB, which describes LLDP.

During the analysis of different implementations on the market, we found several problems of LLDP as well as issues with the interaction to a network management system. These are described in the later chapters.

1. Introduction

The physical topology of a computer network is something every network administrator needs to know. If there is a problem with a connection, a station or a device, the source can be isolated faster.

With devices supporting the Simple Network Management Protocol (SNMP) there are already solutions that support the administration task. Network management software can show a manually entered topology map extended with automatically obtained information from network devices.

This paper describes a method to discover the topology automatically. According to IEEE 802.1AB [1] devices can implement the Link Layer Discovery Protocol (LLDP) and provide more information about the topology over SNMP. The newest specifications for the automation network PROFINET are specifying the implementation of LLDP and SNMP for two of its three device classes.

2. PROFINET

PROFINET is an automation network, based on and compatible with Ethernet (IEEE 802.3) and specified in IEC 61158 [5] and IEC 61784 [6]. A PROFINET IO system consists of an IO-controller, one or more IO-devices and possible IO-supervisors. The IO-supervisors are typically engineering tools.

2.1. IO-system classes

The specification provides three conformance classes of PROFINET IO systems. These classes differ in the supported application-, communication- and redundancy-classes and specify the required features.

Class A specifies certified IO-controllers and IO-devices with standard Ethernet interfaces and standard Ethernet network infrastructure. Class B requires in addition to Class A that the network infrastructure conforms to the PROFINET specification. LLDP and the SNMP are required to support Class B and thus additional management features of the Ethernet network are allowed. In Class C PROFINET IO systems, additional services for redundancy and isochronous protocols are also mandatory.

2.2. Objectives

Different objectives may be reached with the two protocols LLDP and SNMP: automatic topology discovery, automatic device addressing and traffic scheduling.

2.2.1. Automatic topology discovery

It is common for existing fieldbus (e.g. INTERBUS) users to connect a configuration tool to a fieldbus and automatically discover the topology and configuration of the network in the field. It is unnecessary for the planning engineer to enter all this data into the tool. A new fieldbus based on industrial Ethernet must offer at least the same service. With LLDP the information about the topology is stored in every device and SNMP allows the collection of this information to build the topology map.

2.2.2. Automatic device addressing

For management purposes Ethernet devices usually have IP addresses distributed by the IT department. A device itself has a worldwide unique MAC address. When it is replaced, the MAC address changes. The new device in place must have the same IP address as the previous one. So a mechanism for address distribution has to be used.

Normally, PROFINET uses manually distributed names to recognize a device and give it the

corresponding address. However, if a network supports topology discovery, the distribution of addresses may be based on the previously defined topology information. If an IO-device is to be replaced, it must be ensured that it is connected the same way as the previous device. Addresses will be distributed by the IO-controller depending on the location of the IO-device in the network.

2.2.3. Traffic scheduling

In a Class C PROFINET IO system, the scheduling of the cyclic transmission of the frames is planned (isochronous transmission). For this purpose information about the topology of the network is required.

The IRTflex (IRT = Isochronous Real Time) has to be differentiated from the IRTtop. With IRTflex, the scheduled information is sent in the reserved timeslot at best effort, whereas IRTtop uses the complete topology information to plan the scheduling.

The line delay information from the time synchronisation protocol and the delay time on the different links can be collected for the purpose of planning and scheduling. The switch also provides its own information about the delay time, so a complete timing analysis of a PROFINET IO system becomes possible.

3. Topology discovery

In the past, there were many attempts to realize a topology discovery system. Earlier techniques, as described in [7], focused on layer 3 with ICMP messages like those used by traceroute. These were able to discover the network layer topology.

To discover the physical connections, tools and algorithms based on layer 2 - the link layer - have to be used. Former algorithms used primarily the forwarding data stored in manageable switches, the Address Forwarding Table (AFT). If a switch or a bridge supports SNMP, its AFT is stored in a specified place in the mandatory MIB-II (MIB = Management Information Base).

However, all these methods have the disadvantage of misusing data intended for another purpose. First attempts to address the problem directly took place with the introduction of proprietary protocols like Cisco's Discovery Protocol (CDP). Unfortunately they are of no use in a multi-vendor environment and this led to IEEE 802.1AB, the specification for the Link Layer Discovery Protocol and its LLDP-MIB.

3.1. LLDP

The Link Layer Discovery Protocol (LLDP) according to IEEE 802.1AB [1] is the vendor-neutral successor to Cisco's Discovery Protocol. It is very similar to CDP but not compatible.

A device with fully implemented LLDP can send out LLDP Data Units (LLDP DU) on all its physical interfaces. In return it can receive such packets from its neighbours and save the data in its local LLDP-MIB. Both tasks are managed by the LLDP-agent.

The underlying concept is based on the condition that every physical connection-point needs a Media Service Access Point (MSAP) assigned to it. A MSAP has to be unique in the observed network and it consists of two identifiers: the device identifier („Chassis ID“) and the port identifier („Port ID“), which in turn has to be unique per device.

The packets are made up of several Type-Length-Value-blocks (TLV) that contain specified information about the sender. When a device receives a packet, it writes an entry in the LLDP-MIB, linking the included MSAP with the number from the local port that received the packet. This entry represents a direct physical connection, assuming there were no transparent devices between these two MSAPs.

3.2. TLV

The standard IEEE 802.1AB includes sets of information blocks. As mentioned they are organized in a Type-Length-Value-format, which gives it the acronym TLV. This format allows a flexible and extensible frame structure, as these TLV are being attached consecutively (Figure 1).

There are only three mandatory TLV in a valid LLDP DU, and these have to be in the following order after the MAC-header: "Chassis ID", "Port ID" and "Time-to-Live". The transmission of the three mandatory TLV is a prerequisite for the topology discovery with LLDP. These include the MSAP and the "Time-to-Live" (TTL). The TTL is a hint on the temporal validity of the information sent. The receiving agent can use it to determine for how long it has to keep that information stored in a MIB-Entry.

Afterwards, a device can attach optional TLV to provide more information: "Port Description", "System Name", "System Description", "System Capabilities", and "Management Address". With these optional TLV the NMS and the administrator can make use of more information.

In addition IEEE 802.1AB allows the use of

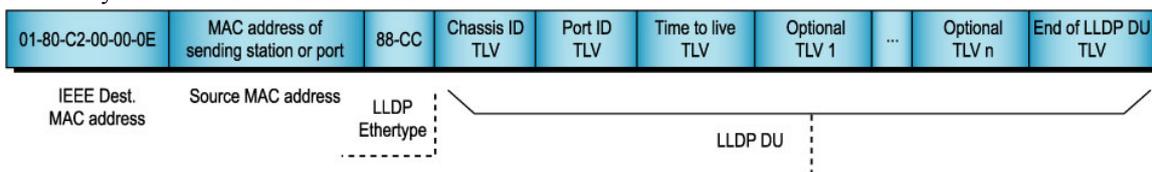


Figure 1: Structure of an LLDP DU [3].

optional vendor- and organisation-specific TLV. However, every new TLV demands an extension of the LLDP-MIB, which will be attached to the “lldpExtensions”-branch. In conjunction with PROFINET an extension was released to exchange information about the port status and delay values. These delay values are necessary for IRT-functionality, where the internal processing time and the transmission time have to be taken into account.

3.3. LLDP-MIB

The LLDP-MIB is the interface between the LLDP-agent and the network management station and can be accessed through SNMP. It consists of several parts: the base-LLDP-MIB and the attached organisation-specific extensions.

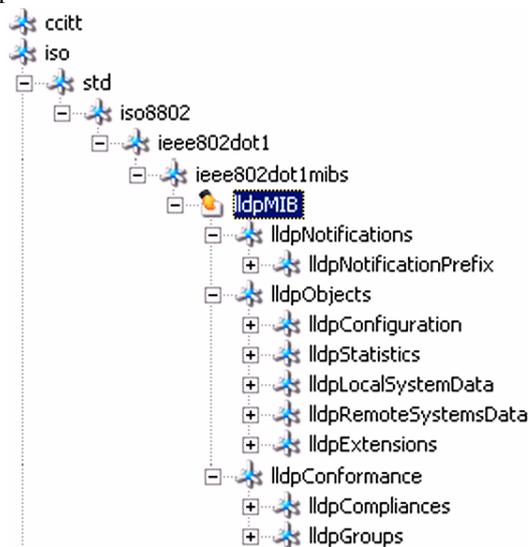


Figure 2: Attaching point of the LLDP-MIB.

The complete SNMP-MIB is structured as a tree and consists of the mandatory MIB-II and other MIBs, which are specified in a standard, RFC or as a private MIB. IEEE 802.1AB describes this LLDP-MIB and specifies the path where it has to be attached in the SNMP-MIB (as shown in Figure 2).

3.4. Network Management Station

The devices exchange LLDP DUs periodically to keep the information about their neighbours up-to-date. However, to discover the actual topology a separate station is required. This station collects all the available data in the devices with SNMP and combines it. The result is a database with all the devices and the connections between them. This can be shown as a topology map.

This station is called the Network Management Station (NMS) and contains the “intelligence” to create a topology map out of the distributed LLDP-data.

With PROFINET, the support is not yet very sophisticated. Some devices do include LLDP-support,

but the associated software does not make much use of it. But the companies promise to include the topology discovery with further developments.

4. Limits and problems

The final version of IEEE 802.1AB was released in May 2005. In the market however it is difficult to find devices with support for LLDP. Out of this limited offer a setup with devices supporting LLDP was tested with a NMS. The goal was to analyse how the topology discovery operates, how effective it is and how far the development has progressed.

Generally, the exchange of LLDP-data between supporting devices worked without problems. During these tests however, there appeared some issues.

4.1. Switches without LLDP-support

All LLDP DUs are sent to the destination address 01:80:C2:00:00:0E, which lies in a specifically defined range. Switches compliant to IEEE 802.1D-2004 (RSTP = Rapid Spanning Tree Protocol, [2]) have to be aware of addresses in this range because packets with this destination address are not meant to be forwarded into other segments.

This is essential because if a packet with this address is received from a switch not aware of this range, it will forward the packet on all ports (like a hub does). The consequence of this behaviour is that all devices connected to the switch will receive LLDP DUs from each other. This in turn causes the devices to think that each one of them is connected directly to all the other devices connected to the switch.

This behaviour could be found on unmanaged SOHO-switches amongst others. They do not support RSTP, so this is not very surprising. Unfortunately there are also newer managed switches which are not aware of the guidelines in IEEE 802.1D-2004. One tested switch even supported LLDP and RSTP, but as soon as the LLDP-function was turned off, LLDP DUs from connected devices were broadcasted.

4.2. Interpretations

The tested NMS-software could not integrate one type of switch because it only supported the mandatory TLV. Other devices used the optional TLV and showed up in the topology map correctly. A reason for this behaviour may be the absence of the “Management Address”. This TLV makes the link between the IP-address (to provide access over SNMP) and the MSAP.

Another problem is related to different interpretations of the specification that showed up in missing port designators. Normally every connection on the topology map is labelled at both ends with the number of the used port. One device however, did not use numbers in the “Port Description”-TLV of its LLDP DU but a short description of its function, like

“unprotected port”. According to RFC 2863 [4] such a text-based description is allowed for the entry “ifDescr”, which is the entry that is recommended to be used for this TLV. The NMS-software was not aware of this possibility and did not show the port descriptions from this device (as shown in Figure 3)

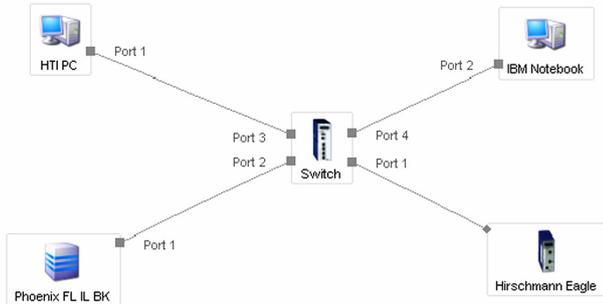


Figure 3: Discovered topology in the NMS.

Some devices also provided wrong information about their own capabilities or they did not use the possibilities of a TLV correctly, e.g. they sent a locally assigned “Chassis ID”, which is not necessarily unique in the network. The use of the MAC-address would be more reliable for instance.

4.3. SNMP

IEEE 802.1AB does not mandate the use of SNMP. But it requires that, if SNMP is not supported, the system should provide “storage and retrieval capability equivalent to the functionality specified in 11.1 for the operating mode being implemented”. Section 11.1 in turn specifies to support the functionality defined in the chapters 8 to 11 with no particular implementation implied.

The result of these definitions is that most of the compliant devices do offer access through SNMP, although this is not entirely certain.

At the moment however, we do not know of any implementation that uses another protocol to access the LLDP data storage. Concerning PROFINET there is a proposal for using the PN-specific protocol DCP. However, with the introduction of the three conformance classes, LLDP support is required only in conjunction with SNMP. Therefore there is no real need for the use of DCP.

It can be assumed that LLDP-devices without SNMP will only play a role in special applications. In this case appropriate NMS-software has to be used, as normal packages only support the data acquisition over SNMP at the moment.

5. Conclusion

This paper describes a new method to discover the physical topology of a network. With devices supporting LLDP and help from a NMS-software it is

possible to handle this task in a comfortable and fast way.

This method is rather extensible (via optional or organization-specific TLV), making it possible to solve problems beyond the topology discovery. The automatic device addressing in PROFINET is one example for an advanced application.

Unfortunately there are some problems still to be solved. Most of them are related to the fact, that the manufacturers are not yet finished with the implementation of LLDP in their devices and software. With more devices on the market, the quality of the LLDP-support should rise proportionally.

Concerning PROFINET, the three conformance classes have to be differentiated. In Class A, where LLDP and SNMP are not mandatory and the use of general network devices is possible, LLDP should be turned off. The mixing of devices with and without LLDP functions could lead to incorrect data and probably add network traffic caused by switches broadcasting LLDP DUs.

In Class B and C the implementation of LLDP and SNMP is mandatory. This preserves compatibility with most NMS-software if all the implementations were done according to the specifications. As no combination of Class B/C PROFINET IO-controller with suitable software was available, we couldn’t make use of the intended functions during our tests.

Particularly for Class C with IRTtop the definitions have to be very precise. Extensive tests with such equipment have yet to be performed.

References

- [1] IEEE: IEEE 802.1AB-2005 - IEEE Standard for Local and metropolitan area networks Station and Media Access Control Connectivity Discovery; available at www.ieee.org
- [2] IEEE: IEEE 802.1D-2004 - IEEE Standard for Local and metropolitan area networks Media Access Control (MAC) Bridges; available at www.ieee.org
- [3] Markus Rentschler: Die Entdeckung der Netzwerk-Topologie, Elektronik 24/2005; available at www.elektroniknet.de
- [4] IETF: RFC 2863 - The Interfaces Group MIB, June 2000, available at www.rfc.net
- [5] IEC: IEC 61158 - Digital data communications for measurement and control - Fieldbus for use in industrial control systems, available at www.iec.ch
- [6] IEC: IEC 61784 - Digital data communication for measurement and control, available at www.iec.ch
- [7] H. C. Lin, S. C. Lai, P. W. Chen: An Algorithm for Automatic Topology Discovery of IP Networks, Proc. of IEEE ICC, 1998, available at www.ieee.org