Network Diagnostics for Industrial Ethernet

Oliver Kleineberg
Hirschmann Automation & Control GmbH
Stuttgarter Straße 45-51
DE-72654 Neckartenzlingen
oliver.kleineberg@hirschmann.de

Max Felser
Bern University of Applied Science
Engineering and Information Technology
Ilcoweg 1, CH-3400 Burgdorf
max.felser@bfh.ch

Abstract

Industrial networks used to control production machinery require high availability to keep possible productivity losses to a minimum. The challenges to overcome are how we are to acquire the needed device information - also with a possible network error already in place - and how to process and pass on this information in order for the error to be quickly fixed. We took possible error scenarios and the data available on the devices to determine which errors can be detected and processed and propose a system of abstraction which provides administrative and service personnel with means to quickly and efficiently identify and fix detected errors.

1. Introduction

More and more, Ethernet based fieldbuses are used in practical applications. There is the promise that Ethernet based solutions will be more flexible and permit the end user a smooth integration of his field equipment in the IT world [1], [2].

Profinet is an automation network based on and compatible to Ethernet (IEEE 802.3) and specified in IEC 61158 and IEC 61784 [3], [4], [5]. A Profinet IO system consists of an IO-Controller, one or several IO-Devices and one or several IO-Supervisors. The IO-Supervisors are typically engineering tools [6], [7].

The specification provides 3 conformance classes of Profinet IO systems. Theses classes differ in the supported application-, communication- and redundancy-classes and are specifying the needed features. Higher classes are compatible to the lower ones.

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 also network infrastructure conformant to the Profinet specification. This means also the support of the LLDP and the SNMP and allows additional management of the Ethernet network. In Class C Profinet IO systems additional services for redundancy and isochronous protocols are also mandatory.

The end user of such an installation is not interested in sophisticated measurement and diagnostics. He needs to know if the installation is correct and error free, if there are some warnings about potential problems or if there are errors in the configuration or installation of the network.

In this paper we list some approaches how to analyze such a network and demonstrate how to integrate such a diagnostic method in a commercial network diagnostic tool for automatic system classification.

2. Measurement and information sources

The first step in creating functional and usable configuration error detection is to acquire the data on which to base the analysis of the topology or the individual device.

Data procurement can be done manually by an operator checking all relevant sources like management interfaces on network devices or capturing network traffic manually e.g. via Ethereal/Wireshark or similar protocol analyzer tools. The obtained information will then be used by the operator to analyze the state of the network and to detect errors within the monitored system.

This method of obtaining information about managed devices is very time and cost intensive, which lead to an early development of several protocols designed to facilitate information procurement from networks and managed devices.

Most protocols and processes in network management are based on the concept of one or few central network management stations (NMS), from which the operator can access all/many managed devices. Generally speaking, a NMS is a focal point of information and control flows within a managed network: It sends out control information, like information requests or configuration change requests to the managed devices while simultaneously receiving data flows in form of data/information packets sent from the managed devices as response to its former requests or triggered by events on the specific device.

One widely used protocol to realize this form of control and information flow between managing and

managed devices is the Simple Network Management Protocol (SNMP), described in its initial version in RFC 1157 [8].

RFC 1155 [9] specifies a “Structure of Management Information” or SMI. The SMI specifies a structural representation of all managed data on a device in which generally every managed object is defined by a name (or object identifier, OID), syntax and encoding rules. Managed objects on devices are often grouped together in software components similar to small databases, called Management Information Bases (MIBs) [10], [11]. This provides the SNMP with means of addressing specific values on a managed device and thus, it provides a general means of accessing management information over networks.

Another aspect of network management and management data transmission is the exchange of information between managed devices. One protocol that accomplishes this functionality is the Link Layer Discovery Protocol (LLDP), a standardized enhancement to the proprietary Cisco Discovery Protocol (CDP), defined in standard IEEE 802.1AB-2005 [12].

The LLDP transmits management data from a device, stored in a corresponding LLDP MIB, in the form of standardized LLDP Protocol Data Units (PDUs), on all network interfaces of a specific device. The LLDPDUs are sent to a reserved LLDP Media Access Control (MAC) Multicast Address, which will not be forwarded by devices compliant to the IEEE 802.1D-2004 standard. This means that management information, sent via LLDP from one device on all its interfaces, will reach only its direct neighbours [13].

Thus, each device receives information about their configuration from its direct neighbours over LLDP and stores these neighbour information in its own LLDP MIB. This means that every LLDP MIB on every supporting network device contains a management data representation of the individual device and all its adjacent devices. This information about each device and its neighbours can now be obtained using SNMP.

Another, completely passive method of information gathering is the possibility to collect data from basic information sources, integrated into the devices’ hardware or firmware itself: statistical counters. Those counters can be integrated into the equipment, e.g. on the level of the PHY integration, to measure the number of network collisions or runts on an Ethernet Network.

The information gathered from these sources can provide an evaluating operator with general information about the state of a network and in case of an error on a network which can’t be pinpointed to a specific device or component, with information to isolate and detect the erroneous device.

3. Detecting misconfigurations

The information acquired through these sources can be used to either detect permanent network/device misconfigurations, or to detect sporadic errors, not visible through comparison of static configuration data alone.

3.1. Detecting a static misconfiguration or error

Configuration data that was exchanged between neighboring devices, e.g. via LLDP, is compared against a defined set of rules which specify erroneous configurations. The set of rules the devices configuration will be tested against are dependant on the network type, the devices that are used, their capabilities, configured protocols and the context of usage, e.g. an industrial Ethernet network. Profinet-IO specifies the requirement to have available at least a 100 MBit/s full-duplex network. This can only be tested by the IO-devices or IO-controller for the link connected directly to the device itself. There is the need to verify this network configuration rule in the whole network also for links between the different switches.

As with LLDP properly configured, each device has information about itself and its immediate neighbor, for every device and every link between individual devices, error detection on the basis of comparing local and remote configuration information can be performed and misconfigurations that could/would lead to errors within the network can be reported to the NMS. The LLDP, in this case, is especially useful for this kind of application as its LLDPDUs traverse network links which are inaccessible or blocked for higher protocols such as IP/SNMP. This ability enables the LLDP to provide management information to neighboring devices even over potentially misconfigured connections between devices.

This can be best explained in a simple example of detecting a static misconfiguration. In this case, on two adjacent devices and their interconnecting network ports, VLANs (Virtual Local Area Networks) [14] have been

configured as shown in figure 1.

Figure 1. Static VLAN misconfiguration.

Each Switch has three VLANs configured on its connecting port, both have VLAN IDs 1 and 2 configured. With the third VLAN ID however, there seems to be a problem. Switch 1 has configured ID 10, while Switch 2 has configured ID 8. This means that

frames tagged with VLAN ID 10 from Switch 1 won't be travelling past Switch 2 and ID 8 tagged frames from Switch 2 won't be transmitted further by Switch 1 as well. This could lead to a loss of connectivity for parts of the network dependant on this VLAN.

This misconfiguration can be immediately detected, providing that both devices implement the use of LLDP, LLDP MIBs and support TLVs (Type Length Value) Fields in the LLDPDUs that transmit VLAN information such as the IEEEs own organizationally specific TLV "Port and Protocol VLAN" TLV described in Standard IEEE 802.1AB [12] .

3.2. Error detection with contextual information

With the same VLAN information present on both devices, the misconfiguration can also be detected on both devices, on each in the devices own context: On Switch 1, the VLAN ID 8 configured on the remote device (Switch 2) could be considered to be falsely configured. In the context of Switch 2 however, VLAN ID 10 on its remote device (in this case Switch 1) could be falsely configured.

This shows that static misconfigurations can be detected by the means of comparing LLDP local and remote information, but to further ascertain how to evaluate and possibly remedy a detected misconfiguration, the context in which it is viewed needs to be broadened up. For this example, the context can be broadened up by checking VLAN configurations on each connected port of the devices and check whether VLAN ID 10 or 8 is configured on further ports and other adjacent devices in the topology. Error detection in the context of each device is limited to the amount of data each individual device has access to and in the case of LLDP, it is limited to itself and all adjacent devices.

There is also the possibility of dynamic configuration error. As soon as the configuration changes, additional error conditions may arise. This is mainly the case if a network infrastructure is modified by the maintenance team: as soon as a reparation operation is executed it may also be the source of misconfiguration. So after every small modification a static configuration test has to be executed.

3.3. Detecting a sporadic error

Sporadic errors can't be detected on the same basis of static configuration errors alone, as a sporadic error is not always present and therefore cannot be effectively traced back to a specific network configuration. Statistical counters can provide a first method of analyzing and identifying sporadic errors. For example, a flapping network link can be identified by a counter set to measure the number of link up and link down operations on a monitored interface. As soon as the sporadic errors become undetectable by counters alone, because of complexity and or the lack of frequent incidence, other methods need to be considered. For sporadic errors, actively trying to reproduce the error and therefore trying to narrow down and detect the parameters responsible for the error is another possibility of detection. This can be achieved by developing and implementing test patterns which change specific single items of a network system's configuration until the error is observed. After observing the error, another parameter is changed and tested and so on. This iterative process is done mostly manually today because of the level of complexity of modern (industrial) networks and the corresponding devices.

4. Error classification and error notification

Another aspect of error detection is how the detected errors are classified and how they are presented to the operator.

The level of complexity of modern networks leads to a very high number of possible misconfigurations and thus possible detectable errors. How those errors are presented to the operator and how the error detection aids the operator in evaluating each error on the basis on how critical it is to the mission, is a crucial part of its usability and has a significant impact how it can and will be generally accepted.

Therefore, first we need to determine how a network error manifests itself to a user or a device using the network for data transportation. From an end users perspective, which can be a normal user, a server or even a production machine, the network itself is a working asset, much like a tool, a pen or a desk.

Therefore, from a user or possibly a production machines "point of view", the network can have five distinct states which have different impacts on its working performance:

1. The network is working normally without any faults or failures. This will result in normal working performance of the user or device. No action of the operator is required.
2. The network is working normally. It was designed to provide media redundancy but there is a fault providing that the redundancy is not available. Repair action is required.
3. The network is not working normally, a loss in performance can be observed while the network service itself is still available. This can result in a reduced productivity by the user or device. Urgent action is required to provide full service.
4. The network service is interrupted and the data transportation services are not available. This can result in a complete loss of productivity of the user or device.
5. No information about the network state is available, as the network is not accessible: it is like a black spot.

All errors in a network configuration can be aggregated to those five states, which are also the

relevant states for an operator making immediate decisions on how to prioritize fixing errors on all participating network devices. Therefore, a misconfiguration leading to the total loss of connectivity has to be prioritized higher than every misconfiguration leading only to a possible loss of performance, as the goal must be to make certain that the total loss of productivity is minimized.

A classification on the five states could be as follows:

- State 1 would be defined as “Network Status OK” or “no fault or failure present.
- State 2 would be defined as “Redundancy fault detected”.
- State 3 would be defined as “Configuration fault detected”. This means that the network is still operational, but network availability could be in jeopardy and network performance may be impacted.
- State 4 would be defined as “Failure detected”. This means that the network is no longer available due to a critical misconfiguration that interrupts the networks operation.
- State 5 would be defined as “Status unknown”, which in this case has the highest criticality. The first and foremost action that is to be taken is to regain reliable information about the state of a network so that further measures can be implemented if necessary.

All states are implicitly sorted in ascending order of criticality. Depending on the scope of the error detection system, each of these classification systems can be implemented into a software application used for network management and/or network surveillance.

5. Summary

With SNMP/MIB, LLDP and dedicated devices/services like statistical counters, a solid technological foundation on how management data is represented and transmitted has already been made available for use in systems for further processing. These mechanisms can now be used to detect errors and misconfigurations on a purely passive basis by comparing management data e.g. transmitted via LLDP or acquired by statistical counters of different devices. When interpreted in a broader context or tested against a set of rules, network errors and misconfigurations can be detected, evaluated and presented to a user or administrator. For the detected and presented error information to be of use for a purposeful and efficient tool for error detection, an abstraction from the high number of possible levels of importance and severity needs to be done. This is achieved with the approach of aggregating errors and misconfigurations to states that are directly mapped to effects on the actual productive network.

This leads to the possibility of creating tools for network administration which detect misconfigurations between devices and errors present in a monitored network with the focus of providing a means of quickly judging situations and based on the findings, taking the right steps to quickly resolve critical situations and remedy faults and failures that lead to a loss of productivity.

References

[1] Felser, M.: “Real-time Ethernet - industry prospective”, *Proceedings of the IEEE*, Volume 93, Issue 6, June 2005, Page(s): 1118 - 1129

[2] Decotignie, J.-D.: “Ethernet-based real-time and industrial communications”, *Proceedings of the IEEE*, Volume 93, Issue 6, June 2005, Page(s): 1102 – 1117

[3] IEC 61784-2: “Industrial communication networks – Profiles - Part 2: Additional fieldbus profiles for real-time networks based on ISO/IEC 8802-3”, available at www.iec.ch

[4] IEC 61158-6-10: “Industrial communication networks – Fieldbus specifications – Part 6–10: Application layer protocol specification – Type 10 elements”, www.iec.ch

[5] IEC 61158-5-10: “Industrial communication networks – Fieldbus specifications – Part 5–10: Application layer service specification – Type 10 elements”, www.iec.ch

[6] J.Feld: PROFINET – Scalable Factory Communication for all Applications, 2004 IEEE International Workshop on Factory Communication Systems, September 22 – 24, 2004, Vienna, Austria, page 33 – 38, www.ieee.org

[7] PROFIBUS International: PROFINET: Technology and Application, System Description, Document number: 4.132, Issue April 2006, available at www.profibus.com

[8] IETF RFC 1157 - “A Simple Network Management Protocol (SNMP)”, May 1990, available at www.rfc.net

[9] IETF RCF 1155 - “Structure and Identification of Management Information for TCP/IP based Internets”, May 1990, available at www.rfc.net

[10] IETF: RFC 2863 - The Interfaces Group MIB, June 2000, available at www.rfc.net

[11] IETF: RFC 3148 - Management Information Base (MIB) for the Simple Network Management Protocol (SNMP), December 2002, available at www.rfc.net

[12] IEEE: IEEE 802.1AB-2005 - IEEE Standard for Local and metropolitan area networks Station and Media Access Control Connectivity Discovery; www.ieee.org

[13] Schafer I., Felser M.: “Topology Discovery in PROFINET”, *12th IEEE Conference on Emerging Technologies and Factory Automation*, 2007. ETFA 2007, September 25-28, Patras, Greece

[14] IEEE: IEEE 802.1D-2004 - IEEE Standard for Local and metropolitan area networks Media Access Control (MAC) Bridges; available at www.ieee.org