Management and monitoring of OT networks, lesson learned from Telcos

Fabrice Strevens, Elvexys SA, Switzerland, Fabrice.Strevens@elvexys.com

Introduction

The aim of this paper is to share best practices discovered by telecommunication operators, as many of their procedures can be of great interest for OT[1] network operation.

Most of the experience has been gathered while working for the historical carrier of Switzerland, as the author was working in the network management, before being involved in the energy industry.

Most of telecoms operators have converged their siloed communication networks to multipurpose IP infrastructure. However, we have seen the emergence of the "out of band" management network, to configure and manage the main data carrying infrastructure. These "out of band" management networks appear to be similar the OT network found in the energy utilities.

Therefore, the kind of operation and management observed at telco can be a great source of inspiration for the OT network operated by energy utilities. In Switzerland, the "out of band" management network of the main telecommunication operator is spread over the country, enabling the NMC[2] to reach and monitor more than 3000 devices located in hundreds of technical rooms.

Organization

Two kinds of staff can be found in such organization:

- **The networking engineers**, who are responsible for the network architecture design, the infrastructure extension, like adding switches, routers, firewall. Finally, they are in charge of the troubleshooting when errors are discovered.
- **The Network Management Center** (NMC) personnel, whose role can be compared to the "dispatch" function in the context of energy grid operation. The duties of the NMC lie in:
    - Supporting connection of new hosts by opening ethernet ports on switches, testing connectivity, and updating the network documentation.
    - Monitoring the device connectivity and analyzing the activity logs send by the network devices, to understand problems, and take proactive corrective measures.

In brief, the network engineering team is responsible for the design, implementation and troubleshooting of the network, while the NMC personnel ensure the daily operation and monitoring.

[1] Operation Technology networks are the field IT infrastructure dedicated to substation operation.
[2] Network Management Center

Documentation

Proper network operation relies on an accurate network documentation. The documentation can be split according to the internetwork layer or the simplified OSI model.

- **Infrastructure schematic (L1)**: these schematics should indicate the wiring between the network elements (switches, routers, firewall). The switch port number and the building patch panel port number should be indicated on the schematic, or if too complex, this information would be reported on a separate spread sheet.

  The infrastructure schematic should provide to the NMC with an accurate view of the interconnection between the networking devices, the NMC should be able to use these documents to remote assist the field technicians while locating and testing the interconnection.

- **Vlan schematic (L2/L3)**: these schematics should show the distribution of the different vlans, subnets and display the default GW. They are be used by the networking engineers for troubleshooting the IP transport.
- **FW Zoning schematic (L4)**: these schematics should show how the subnets are logically grouped in security zones. The rules between security zones are then listed in a dedicated spreadsheet.
- **Host inventory**: Finally, a spread sheet should inventory all the existing IP hosts. For each host, the following information should be available
  - Hostname, also referenced in a DNS record
  - Device description (if the hostname is not self-explanatory)
  - Installed firmware
  - Network interfaces, (the below information shall be given for each network interface)
    - Name
    - Mac address
    - Connected to
    - IP addresses & vlans
    - Media type
    - Interface configuration (10/100/1000/auto, full/half/auto duplex)
  - Physical location (City, Building, Room, Row, Rack, Heigh)
  - Contact team / person

The host inventory may appear time-consuming to establish, however this information represents an excellent snapshot of the connected devices. First, this helps the NMC to remote support a field engineer to physically find a faulty device thanks to the location information. Then the connection information is used by the networking engineer for troubleshooting.

Monitoring practices and associated tools

As the network traffic within IT and telecommunication infrastructures are intangible and invisible, the purpose of the monitoring tools is to provide the operators with a visualization and metrics to understand what really happen into the network. Below is a list of tools and associated protocols sorted in order of importance.

ICMP

One of the most used protocols to monitor the connectivity of the IP hosts remains ICMP, Internet Control Message Protocol. Ping is the well-known trouble shooting application using ICMP. The monitoring server sends an echo request at a regular intervals (i.e. 1 minutes) to the hosts, which answers immediately with an echo reply. The ping tools also provides information like the round trip delay of the echo requests and the number of routers crossed. This last parameter is expressed by the Time To Live (TTL), which is a counter starting at 128, and be decreased by each router crossed on the route to the targeted host. If the TTL reach zero, the packet is discarded as it would mean that the packet is forwarded along an infinite (looping) route.

All monitoring tools include the visualization of the ICMP, as this represents the basic immediate connectivity status and answer to the basic question: “Are the hosts connected and reachable?”.

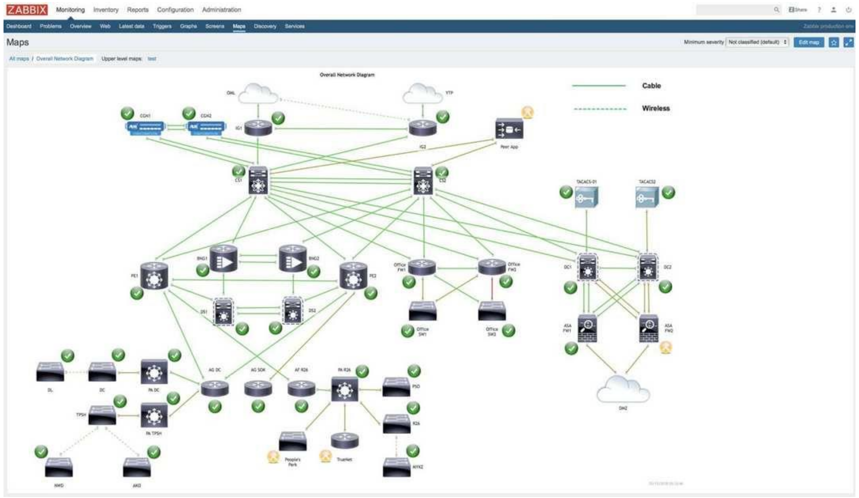

Figure 1 Example of a map view established with Zabbix tool[i]

From ICMP measurement the following information can be extracted.

Round trip delay stability over time, the round trip delay should be stable over a long period of time, ie. Months. A change is the average round trip delay may indicate a change into the structure or in the load of the network. A dramatic degradation, i.e. 100% increase of the average round trip delay, should attract the attention of the NMC and lead to further investigation to understand the reason of the change.

Short term stability of the roundtrip, aka jitter, could be monitored. The latency should stable between echo requests. A variation of the round trip delay, (i.e. 4 ms, 10 ms, 4 ms, 95 ms, 5ms), should be investigated first to understand the reason and to ensure stability of the data packet transport.

Time To Live, the value of the Time To Live should be identical during a ping session (in command line: ping HOST -t, HOST the targeted IP host). A variation of the TTL indicates that the echo packet sent to the host is taking different routes between the requests. This might indicate an issue with the routing configuration of the network. A change of the TTL value over a longer period of time, may indicate a change in network topology. This in itself is not a problem, except if this change cannot be explained by the NMC.

SNMP

The Simple Network Management Protocol (SNMP) allows a network management station, usually in the vicinity of the NMC, to query monitoring information from a network host by polling. An advantage of SNMP is that the parameters to be monitored can be queried via their OID (Object Identifier). The OID provides standardized access to the parameters across the vendors.

The values retrieved via SNMP can be easily graphed with monitoring tools (i.e. the open-source cacti www.cacti.net). These time series charts are useful when NMC or Network engineer would look for correlation between events.

Finally, alert on threshold can be set, which is an important feature when thousands of devices must be monitored, and analysis of the individual equipment are no longer possible.

Fundamental parameters to monitor

As the hosts expose a large amounts of information, it appears that the most interesting parameters to monitor are the following.

CPU Load: the CPU load represents the activity of the device, usually the CPU load of a PLC or a server should be constant and low. The CPU load indicates that the host is actively processing data. A drastic change in the CPU load behavior should be investigated as it suggests that something is happening on the device. It may also be the consequence of a change in configuration, with the device itself or in another network device, which will involve higher workload on the monitored node. The latter can be observed when the routing rules have been updated and are spread over the entire network.

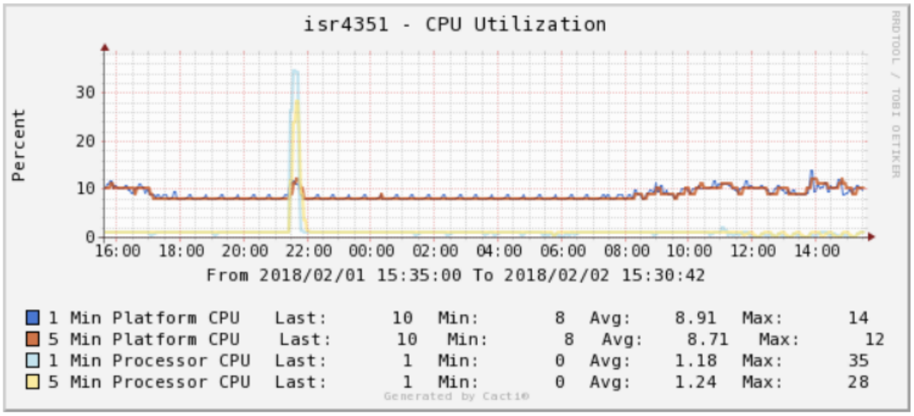

Figure 2 Typical CPU usage of a router[ii]

Memory usage, the device memory usage shows the available RAM for operation. This parameter provides an interesting insight as it shows the memory used by the different applications inside the host. Also, the memory usage should remain constant, or may be in a sawtooth shape, but a continuously increasing average memory usage will unavoidably lead to slow operations or device crash. This phenomenon is called memory leak and come from a bug in the software. This case is quite complex to detect because a memory leak may take years to appear and may occur only with specific configuration. Therefore, watching this parameter will help the NMC and engineer to better understand the internal behavior of the device under monitoring and react before running out of memory.

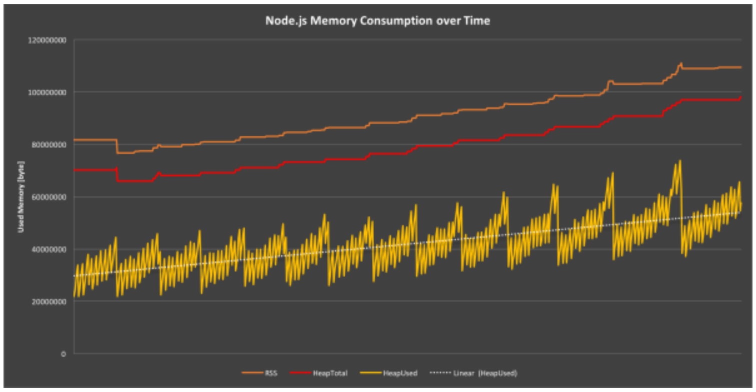

Figure 3 Example of a memory leak[iii]

Disk usage, although less interesting on IED, is of great importance on servers and network devices as router / switches. The disk usage, as on a laptop computer, indicates the available storage. It should be constant, but without any automatic cleanup process, the disk or the nonvolatile memory, over the course of a year will be filled with logs or debug information. Ultimately the device will stop operating when running out of disk space. This may seem obvious, but a device put under operation for 15 years, needs to have this simple monitoring to ensure long term operation.

Device temperature, this parameter simply indicates the level of dust on the systems under monitoring. When the host is equipped with fans this can evolve quickly over a few years. On fanless system, as found in energy substation, the performance degradation of the cooling is less observable.

Advanced parameters

The four parameters above are essential to ensure a continuous long term (decades) operation. Further parameters can be monitored to ensure more a detailed understanding of the network health status.

Network bandwidth: SNMP allows all the network interfaces to be monitored against the used bandwidth. This provides information on available bandwidth for further traffic evolution and thus to avoid future congestions. Also, this is the only available tool to understand the propagation of network burst. On a highly predictable infrastructure like energy Utilities, a burst can occur when firmware is uploaded, or when COMTRADE files are retrieved after a disturbance. As network administrators start to use network bandwidth to monitor the traffic, more detailed information about the traffic might be desired, and sFlow would be a great help to understand the nature of the network traffic (See chapter below).

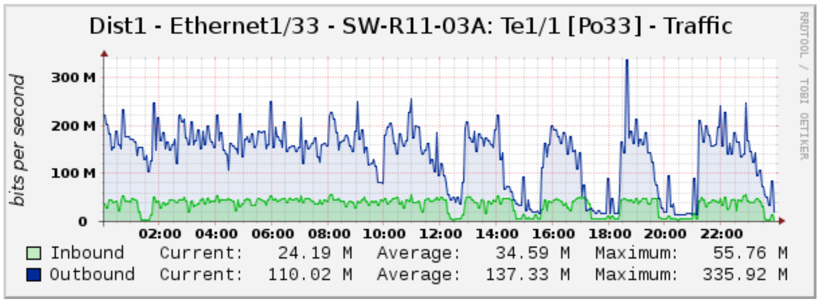

Figure 4 Bandwidth monitoring example[iv]

SFP reception power: The optical power reception can be interesting to monitor as it will underline an aging SFP transmitters. Moreover, any alteration of the optical path between the device will appear on the monitoring chart. A misconnected optical fiber, an aging SFP, or the insertion of a passive monitoring system (Optical TAP) can be detected with this monitoring.

Finally, the two last parameters can be monitored eighter via SNMP or via SYSLOG. The benefit of SNMP is the convenience for graphic generation, even if the chart shows almost no variation.

Reboot or uptime: a restarting device is an abnormal event on the network and must be detected. If a watch dog triggers the restart of the device, this indicates a hardware malfunction of the system. Without monitoring the device continues intermittent rebooting untill a permanent failure occurs. If the device is intentionally rebooted by a human, the root cause is understood, as rebooting is often a way to change a firmware or to apply a new configuration. Monitoring device reboot may help point to suspicious activities like hacker intrusion.

Firmware version or Firmware hash: any unannounced change on a device firmware should be treated as a potential cyber security concern. Usually cybersecurity attacks are “slowly” built by taking control of the devices connected to the network. From an intruder point of view, deploying a corrupted firmware is a way to ensure that a device still continues its normal operation but is ready to be used by hacker for a larger distributed attack involving many devices.

Syslog

Syslog is used to centralize the system logs of the different devices on the network. The messages are sent by the hosts to the syslog server when an event occur. There is no polling mechanism in syslog. The log messages contain information on the application (called facility) running on the device generating the message and the severity of the sent message (from emergency to debugging information). As the syslog message content are not standardized, a human reading is mandatory, to understand the information which limits the automated handling of this information. Therefore, when possible, the SNMP Trap mechanism should be used instead.

The analysis of the syslog messages highlights any loss of redundancy and in certain cases provides an early warning of security issues.

Loss of redundancy

As the redundant system, like dual power supply equipment, continues normal operation when a failure occurs, therefore monitoring of the redundant systems is mandatory to detect such failures. The principal redundant elements to monitor are as follows:

Power supplies are well known to be the weakest items of the entire IT / OT infrastructure. This is why they are always redundant and hot swappable.

RAID storage: redundant hard drive system as RAID 1 or 5 is to be monitored.

PTP grand master clock change: the redundant PTP (timing infrastructure) should be monitored to detect PTP grand master failover.

Network topology change: this message, issued by the network switches, indicates that a network link as failed, this may be the optical fiber, the SFP or a change into configuration of an adjacent devices i.e. change in RSTP parameters. Investigation to understand the nature of the topology change should be taken.

These last two messages may lead to intensive investigations to understand the reason of the loss of redundancy. However, it is worth taking the time to understand the problem because the next failure the network may become entirely instable and unavailable.

Security issues

Syslog messaging can be a first barrier to detect abnormal security behavior. The following messages should be monitored and to trigger further investigations.

Unsuccessful login, any host should be configured to emit a Syslog message when a login tentative, successful, or rejected, occurs. Multiple login failure can clearly be seen by the NMC. These login failures may be a signal for a hacking attempt.

Reboot of a host should be considered as suspicious. The NMC should be trained to sort out the cases, a reboot occurring in the middle of the night associated with a firmware version change should ring all the alarm bells, while a single device reboot while field engineers are announced to be on site may not be considered a concern.

Connectivity error with 802.1x or sticky mac, on a stable network environment like a power substation, enabling 802.1x port-based authentication would be a great security improvement to avoid the insertion of a non-desired and / or unannounced network devices. Notice that sticky-mac mechanism may play a similar role, as it allows only the first connected hosts on a switch-port to be authorized.

While 802.1x or sticky-mac mechanism are activated any tentative to connect foreign devices are monitored via Syslog.

Other issues

Finally, Syslog is of great help to be informed of other issues as port flapping (ethernet port going up-down with no reason). As the syslog messages differ from one host to the other, careful testing and documentation of the error messages should be made in a lab environment prior bringing the system into live operation.

sFlow

The sFlow, stand for "sample flow" and is a lightweight mechanism to gain visibility on the traffic carried by the switches[3]. Indeed, the switches or the probes are generating statistics about the traffic transiting through the switch. These statistics contain, the source, destination IP address, the TCP or UDP port used and finally the number of packets seen.

This information provides deep insights on the network traffic occurring into a substation or along the network. Even if this is not as powerful as a well configured and well understood IDS[4], this sFlow approach is a first barrier to unknown and potential hazardous traffic.

Abnormal traffic like port scans from would be easily identified.

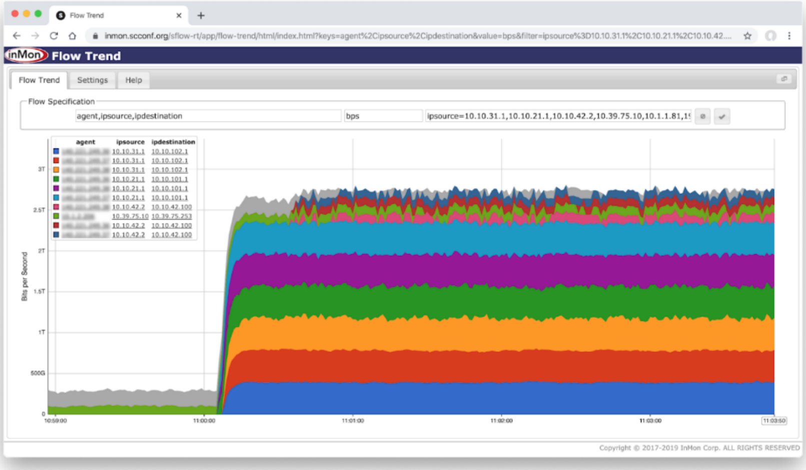

inmon.scconf.org/sflow-rt/app/flow-trend/html/index.html?keys=agent%2Cipsource%2Cipdestination&value=bps&filter=ipsource%3D10.10.31.1%2C10.10.21.1%2C10.10.42....

Flow Specification

agent,ipsource,ipdestination | bps | ipsource=10.10.31.1,10.10.21.1,10.10.42.2,10.39.75.10,10.1.1.81,1!

	agent	ipsource	ipdestination
🟦	[redacted]	10.10.31.1	10.10.102.1
🟥	[redacted]	10.10.31.1	10.10.102.1
🟧	[redacted]	10.10.31.1	10.10.102.1
🟩	[redacted]	10.10.21.1	10.10.101.1
🟪	[redacted]	10.10.21.1	10.10.101.1
🟦	[redacted]	10.10.21.1	10.10.101.1
🟥	[redacted]	10.10.42.2	10.10.42.100
🟩	[redacted]	10.39.75.10	10.39.75.253
🟥	[redacted]	10.10.42.2	10.10.42.100
🟦	[redacted]	10.10.42.2	10.10.42.100

Copyright © 2017-2019 InMon Corp. ALL RIGHTS RESERVED

[3] Certain switch can natively export statistic information about the traffic. As this feature might not be available on substation switches, a network probes connected behind a mirroring port (SPAN Port) would do the job too.

[4] Intrusion Detection System

Figure 5 sFlow provide detail on network traffic[v]

Conclusions

We are now in the era of digital energy substations. As we are at the crossroad of Telecommunications, Electricity, and Software (with XML describing the systems), many engineers feel uncomfortable with the other domains. Looking at the best practices found in other industries could be insightful.

Regarding the Telecoms component, 1st documentation is a key to success. Next, monitoring and understanding the system behavior under normal operation is the second requirement to ensure readiness to react in case of abnormal situation.

[i] Credit : https://bestmonitoringtools.com/create-zabbix-maps-with-examples/, retrieved in July 2021

[ii] Credit : https://layer77.net, retrieved in July 2021

[iii] Credit : https://www.dynatrace.com/news/blog/understanding-garbage-collection-and-hunting-memory-leaks-in-node-js/, retrieved in July 2021

[iv] Credit: https://paulgporter.net/tag/rrd-tool/, retrieved in July 2021

[v] Credit : https://blog.sflow.com/2019/, retrieved in July 2021