Erik Janhunen

BENEFITS OF AUTOMATION IN NETWORK MONITORING

Bachelor of Engineering
Information Technology

2021

South-Eastern Finland
University of Applied Sciences

South-Eastern Finland University of Applied Sciences

Author	Degree	Time
Erik Janhunen	Bachelor of Engineering	May 2021

Thesis title

Benefits of automation in network monitoring — 31 pages

Commissioned by

South-Eastern Finland University of Applied Sciences

Supervisor

Matti Juutilainen

Abstract

The thesis objective was to inspect benefits of automation in network monitoring and ticket handling. The monitoring company had used manual methods to handle network monitoring. This required the company to always assign one employee from the active shift to monitor detected events and create tickets from them.

The company planned to increase production quality by introducing automation which would handle network monitoring entirely. This would free the monitoring employee from their position, allowing the whole shift to focus on ticket processing. The automation was also programmed to automatically close tickets for events that had been cleared to remove unnecessary tickets from the work queue.

The thesis will briefly go through network monitoring as a service and the technologies used with it. In the end, a comparison was done between old and new production environments where incident ticket amounts were gathered from a customer that was transferred to the new production. This comparison would give insight of how automation would improve ticket processing by lowering the amount of tickets that require investigation by employees.

However, data gathered for comparison from the new production did not match what was supposed to be monitored from the customer network. There were issues with monitoring configurations, which made false positives to be ticketed by automation. At the time when these issues were fixed, there were large amounts of tickets created from false events, which had been closed by automation after monitoring was corrected.

The thesis failed to receive proper comparison data for the automation effectiveness. For future studies, it is recommended to use a longer timespan for comparisons, and to include other metrics to measure production quality. These could be accomplished by creating additional comparisons or surveys to receive feedback about the production environment.

Keywords

monitoring, network, automation

CONTENTS

LIST OF FIGURES

LIST OF TERMS

Event	Alerts from networking devices
ICMP	Internet Control Management Protocol
ITIL	IT Infrastructure Library
ITOC	IT Operations Center
ITSM	IT Service Management
LAN	Local Area Network
MIB	Management Information Base
OID	Object Identifier
SLA	Service Level Agreement
SNMP	Simple Network Management Protocol
SNMP agent	Software that sends data from device to an SNMP manager
SNMP manager	Server that receives and processes SNMP events from SNMP agents

TCP	Transmission Control Protocol
TLS	Transport Layer Security
UDP	User Datagram Protocol

1 INTRODUCTION

The commissioning company for this thesis is an IT company that offers network monitoring as one of its services to medium and large-sized businesses. The monitoring is implemented in a customer management network depending on the capabilities of customer networking equipment. The monitoring is not restricted to a specific manufacturer, which allows the customer to decide their own devices for their network.

The monitoring company uses an ITIL service management framework as its core to build services for their customers. ITIL provides common practices for IT companies to build services and guides them towards continuous service improvement. These practices are not necessary to be implement as is, instead the company may modify them to fit their business operations.

“Not every practice in ITIL can be considered ‘best practice’, and for good reason. For many, a blend of common, good and best practices are what give meaning and achievability to ITSM.” (The Official Introduction to the ITIL Service Lifecycle 2007, 4.)

“ITIL is intentionally composed of a common sense approach to service management – do what works” (The Official Introduction to the ITIL Service Lifecycle 2007, 3). As such, the framework offers best practices which could work for the company but are not necessary to be implemented as described in the framework. The commissioning company have adjusted these practices to create a service for network monitoring, which are also guiding the company to upgrade their monitoring service. This have led the company to create a completely new production environment with upgraded software, which will also support automated monitoring.

The network monitoring itself stays the same for the customer devices. The customer contracts define what is monitored and how the monitoring company should react to events received from monitored devices. Some of the events are only included in reports, while those which affect customer production are

ticketed for further processing. These tickets are created as incidents and are processed by employees according to customer contracts. Incidents are unplanned interruptions in customer production, as described in ITIL Service Operation (2017, 46): “An unplanned interruption to an IT service or reduction in the quality of an IT service. Failure of a configuration item that has not yet impacted service is also an incident, for example failure of one disk from a mirror set.”

The monitoring company is responsible for creating these incidents and resolving them in agreed SLA. The company is constantly reviewing customer feedback and tries to adjust their monitoring to provide faster response to their customers. Earlier monitoring was done manually by having one employee constantly monitor received events and creating incidents. The old software was also nearing its end of life, and the little automation what was included still required employee interaction.

The new production with an upgraded software is built with automation in mind, which could handle event detection and ticketing on its own. The automation is programmed to close tickets automatically for events that had been cleared, therefore removing unnecessary tickets from the work queue. Other quality of life improvements were also included that would make ticket processing easier for employees.

2 NETWORK MONITORING AS A SERVICE

As businesses have transferred toward digitalization, more services require network connectivity. The businesses require their services to be available where downtime may have drastic impact on business operations with the possibility of lost revenue. The businesses need to monitor their services and be swift in fixing problems which might occur.

This would require businesses to create departments and infrastructure to host monitoring services and analyze network traffic, but the cost of software, infrastructure, resources and training could be considered too high by the company to build themselves. For these purposes, the company may outsource the monitoring service to another company. ITIL describes outsourcing as a formal agreement to use another company's services: “This approach utilizes the resources of an external organization or organizations in a formal arrangement to provide a well-defined portion of a service’s design, development, maintenance, operations and/or support” (Service design 2007, 75).

According to Haimi & Huovinen (2018, 19) outsourcing and finding cost effective solutions have been booming with companies starting from mid 1990s: “The mid 1990s to the 2010s were the era of outsourcing. The first large-scale outsourcing deals focused on IT infrastructure and support, later extending towards the development and management of applications. The concepts of SaaS, IaaS and PaaS were also born in the 2000s. – – In the era of outsourcing, the primary focus of Service Management was on cost-efficiency, standardization, vendor management and SLAs.”

By outsourcing, the company is investing on the service itself. The company hosting the service is responsible for the service infrastructure and upkeep, which may occasionally require improvements and upgrades. These services may also be offered to other customers, which will allow the company to fully utilize its services.

As for network monitoring, ITIL practices give common guidelines on how the service could be constructed. The commissioning company mainly uses four processes from ITIL: event, incident, problem and change management. These are handled under ITOC service operations department, which are responsible for processing detected events from customer networks and finding solutions for them. All of these processes may have their own SLAs and actions for each customer, which are specified in customer contracts.

2.1 Event management

Event management consist of handling detected events from the monitored devices. Events are first filtered by a monitoring software, and those which are impacting customer production will be ticketed as incidents. Event severity and device priority determine ticket priority and actions to be taken to handle the incident. In figure 1 event lifecycle is shown from detection to event closing. The event is filtered, and if considered significant or the event impacts production, it will be ticketed as an incident.

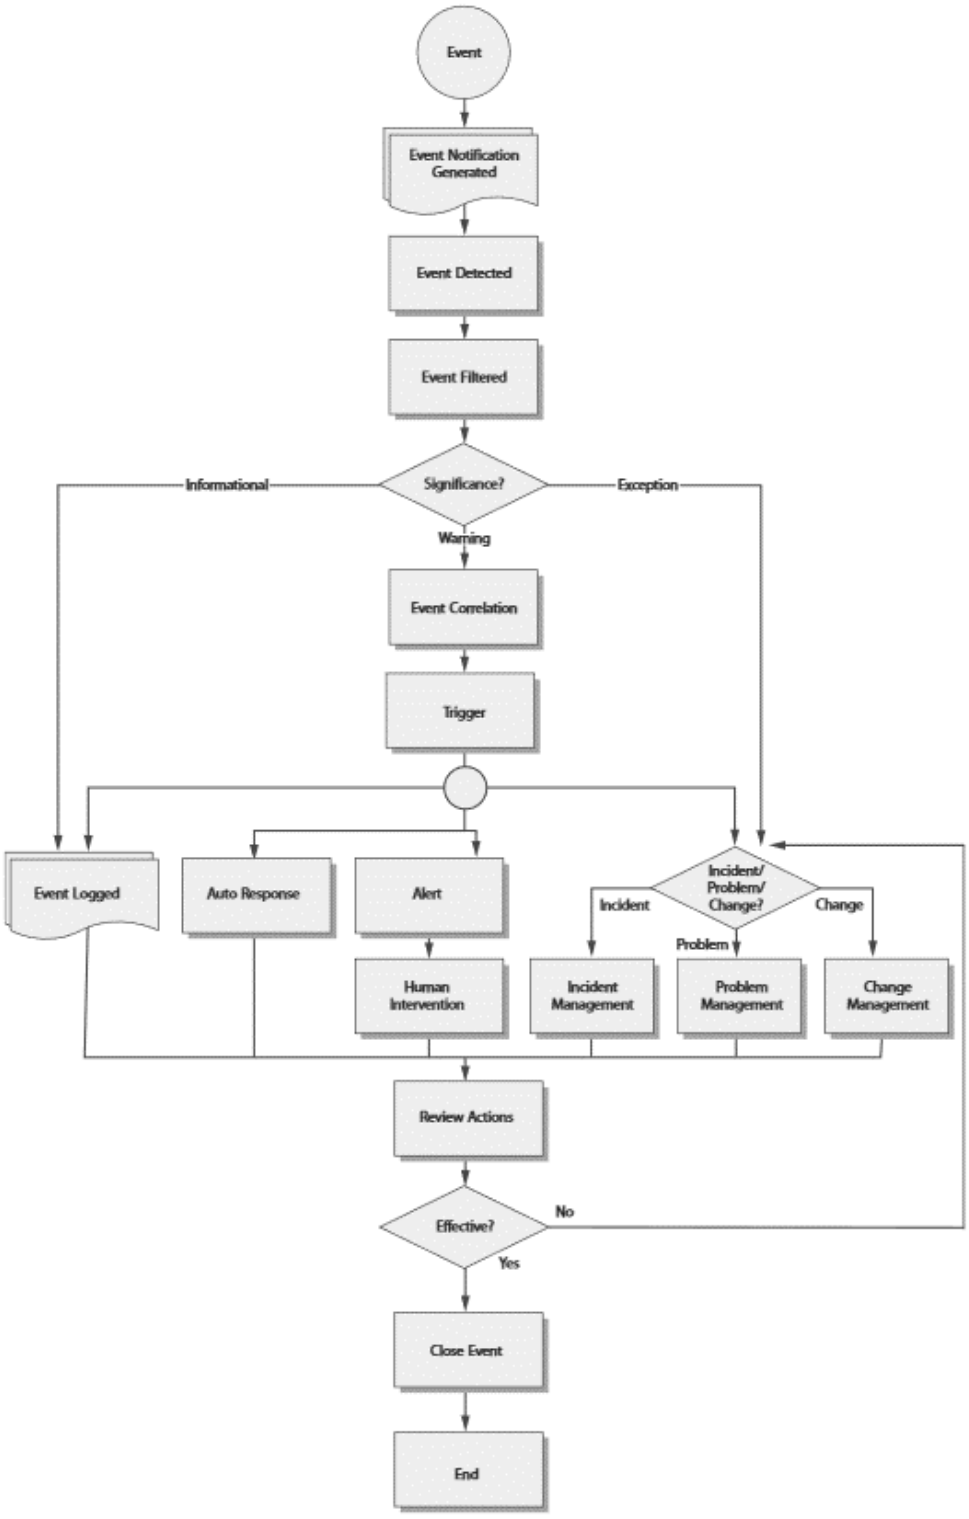

Figure 1 - Event management process flow (Service Operation 2007, 38)

Event processing ends once an effective solution is found. This is found by processing the event through incident, problem or change management. If the applied solution is not considered effective, the event is shifted back to problem management for further investigation.

2.2 Problem management

Problem management will investigate the root cause and a solution for detected event as shown in figure 2. Once the root cause can be identified and a possible solution found, it is added to known error database. If the possible solution requires large changes in the customer environment, the solution will go through change management to identify possible risks and effects which the solution would cause. Once these have been identified the solution will be implemented in the customer environment.

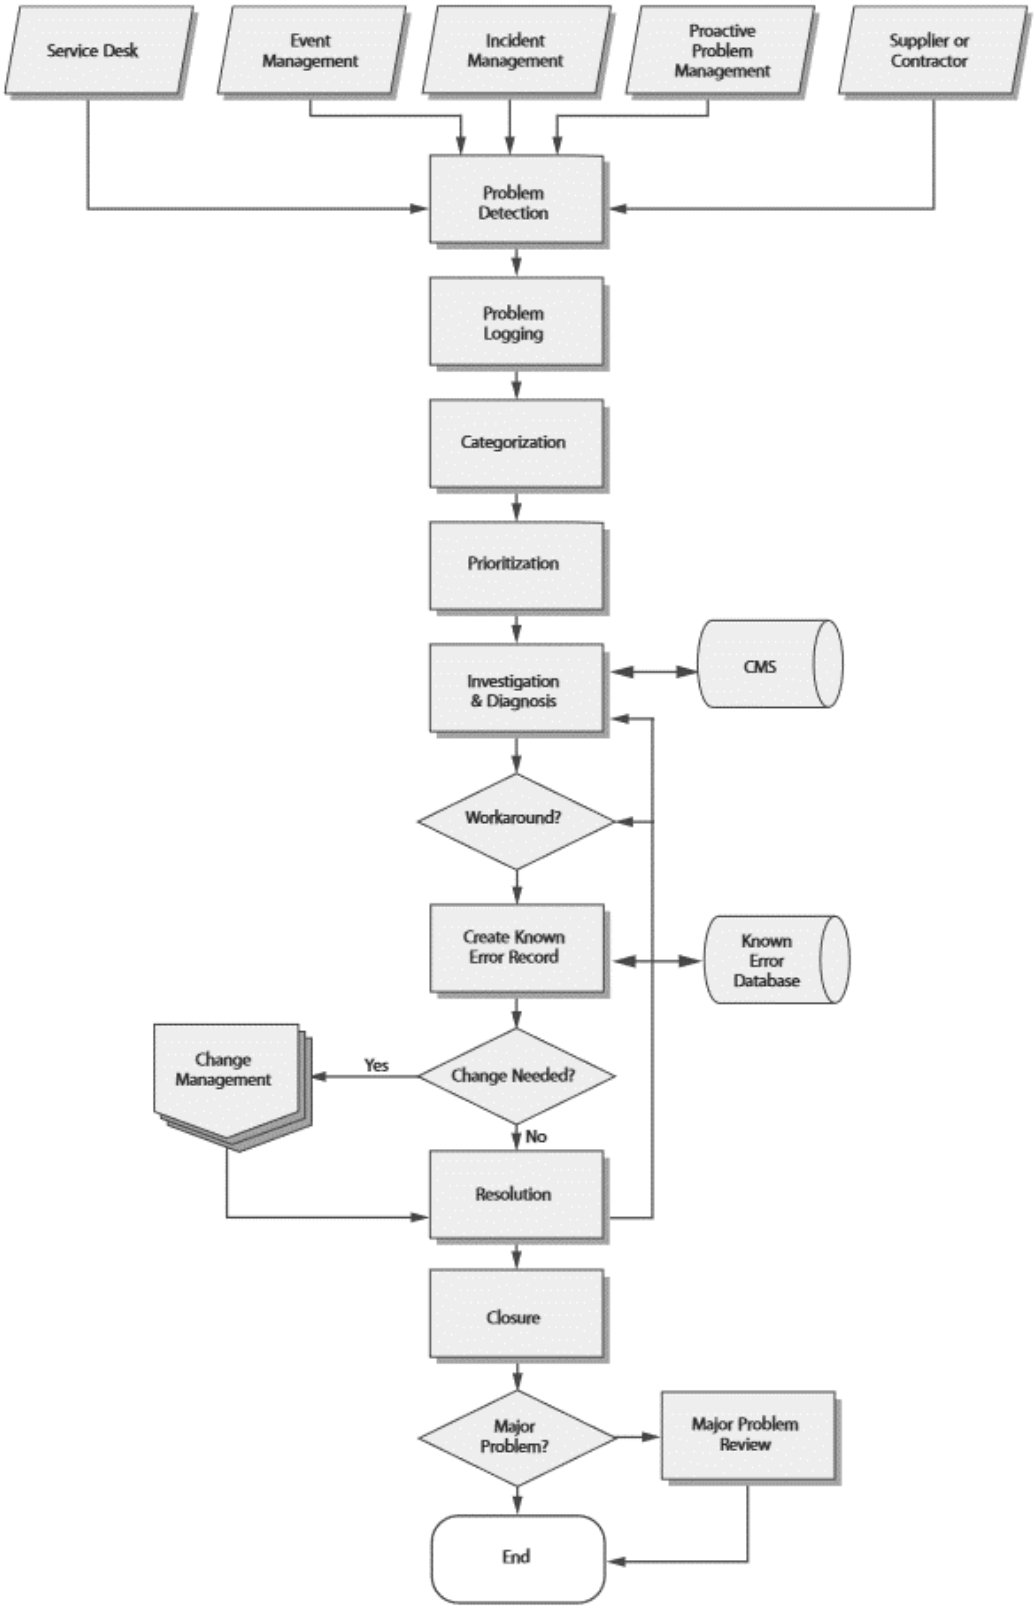

Figure 2 - Problem management process flow (Service operation 2007, 60)

Problem management should also investigate customer networks proactively for possible issues which could prevent the network from failing. ITIL describes the primary objectives of problem management in the following way: “The primary objectives of Problem Management are to prevent problems and resulting incidents from happening, to eliminate recurring incidents and to minimize the impact of incidents that cannot be prevented” (Service Operation 2007, 58).

Therefore, problem management has an important role in production. If the possible issues could be proactively prevented, it would ease work with event and incident management. However, it cannot be accurately predicted when a device might fail, which leaves problem management to create precautions for these failures. The precautions could then be implemented by incident management as a workaround to restore customer production.

2.3 Incident management

Incident management is best described with ITIL: “Incident Management is the process for dealing with all incidents; this can include failures, questions or queries reported by the users (usually via a telephone call to the Service Desk), by technical staff, or automatically detected and reported by event monitoring tools.

The primary goal of the Incident Management process is to restore normal service operation as quickly as possible and minimize the adverse impact on business operations, thus ensuring that the best possible levels of service quality and availability are maintained.” (Service Operation 2007, 46.)

Therefore incidents are not only generated from event management. It is possible that the customer may have detected errors in their production, which they will forward towards the monitoring company incident management as shown in figure 3. If these reports fill the characteristics of an incident, it will be logged to be processed by incident management.

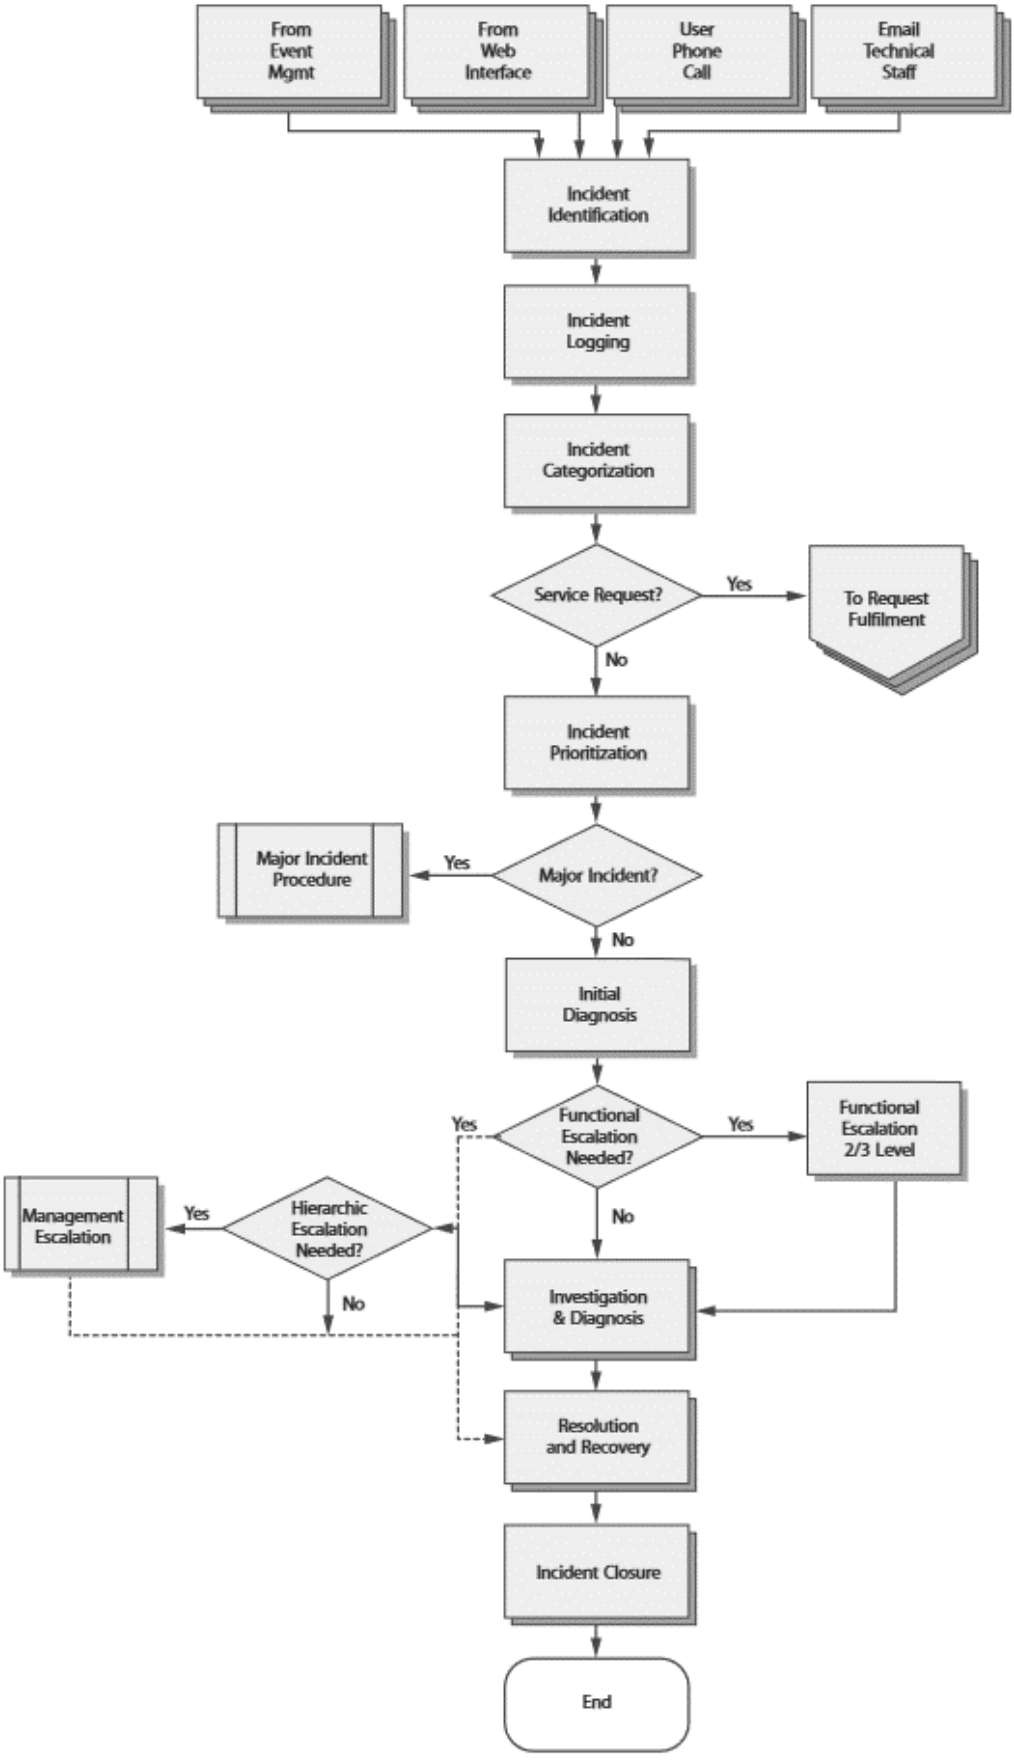

Figure 3 - Incident management process flow (Service Operation 2007, 48)

Customers may also request minor changes to their networks. These requests do not have a high impact on the customer production and may not pose any risks if applied directly. These requests are forwarded to request fulfillment which handle minor changes to customer networks.

Incident management operations may conflict with problem management. If thorough investigations would need to be done, the event should be active or enough knowledge should be gathered to reproduce the problem in a lab environment. Solving an incident without gathering data might prevent further analysis for the issue, which could prevent a permanent solution from being found. However, returning the customer production to normal should always be the top priority.

2.4 Network monitoring

The customer network monitoring is handled by a monitoring software. This software is connected to the customer management network, from which it monitors and receives events from the devices as show in figure 4. Each customer is recommended to segment their network and to have a dedicated management network to further improve network security.

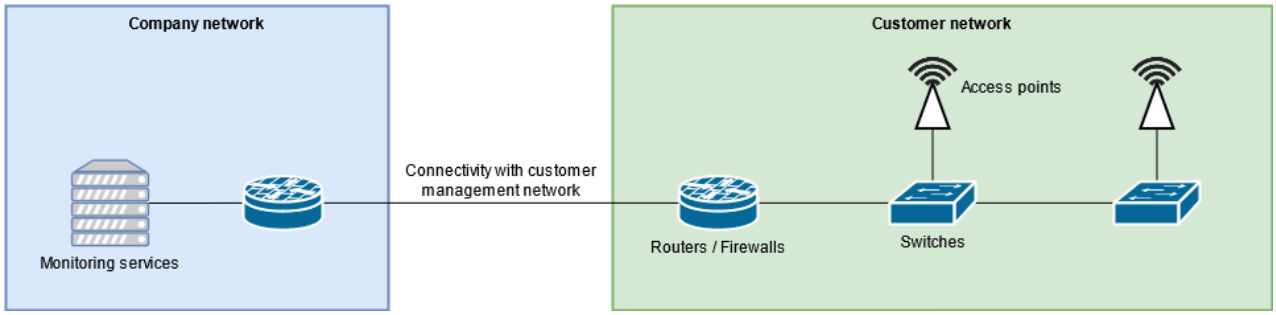

Figure 4 - Connectivity with monitoring company and customer

The monitoring protocols mainly remain the same, where only configuration is required for different manufacturers' devices. Each customer has their own requirements for network monitoring, so configurations and what is being monitored will be different.

3 MONITORING TECHNOLOGIES

Network monitoring is mainly done by using ICMP, SNMP and syslog monitoring. The monitoring software is configured to query and receive data from the networking equipment, and only the required components are being monitored, which are negotiated in the customer contracts.

3.1 ICMP

ICMP was defined by standard RFC792 in 1981. In network monitoring this protocol is used with a ping tool. The tool sends Echo Requests to the host IP address, and the targeted host responds with Echo Replies as shown in figure 5. If the reply is not received from the host, the host is determined to be unreachable.

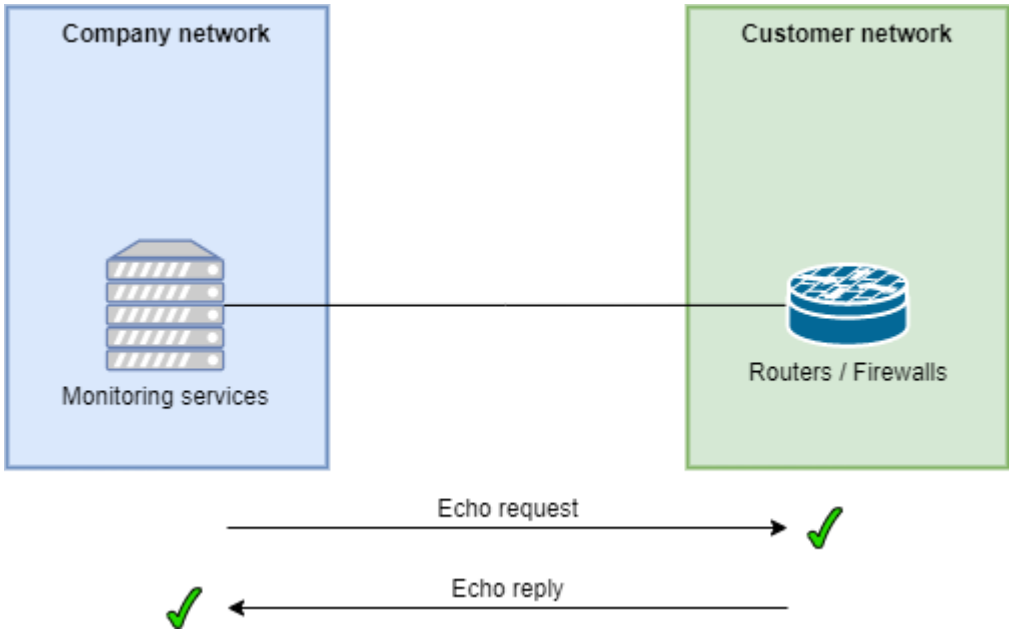

Figure 5 - ICMP monitoring example. Monitoring software sends an echo request to the monitored network device

This monitoring can be implemented on all network devices configured with an IP address, but it does not give any other information than whether the targeted device is reachable. To further troubleshoot unreachable hosts it is required to verify if the network is reachable, and then further troubleshoot the LAN. There could be a connectivity issue on the site between network devices, the device

configuration may not allow answers for echo requests, or the device itself could malfunction preventing it from processing any network traffic. All of these problems may require onsite personnel to visit the site and provide console access to the device, if the device is not accessible remotely.

3.2 SNMP

SNMP monitoring consist of agents, managers, and the SNMP protocol as stated in RFC3411 (2002, 5). The SNMP agents are software on networking devices which send information form the device to the SNMP manager. The managers are monitoring software which process received data from SNMP agents. The SNMP manager may query information from agents to verify the component status as shown in figure 6. In the query the manager requests interface status from one of the interfaces on the router, which the router reports to be up.

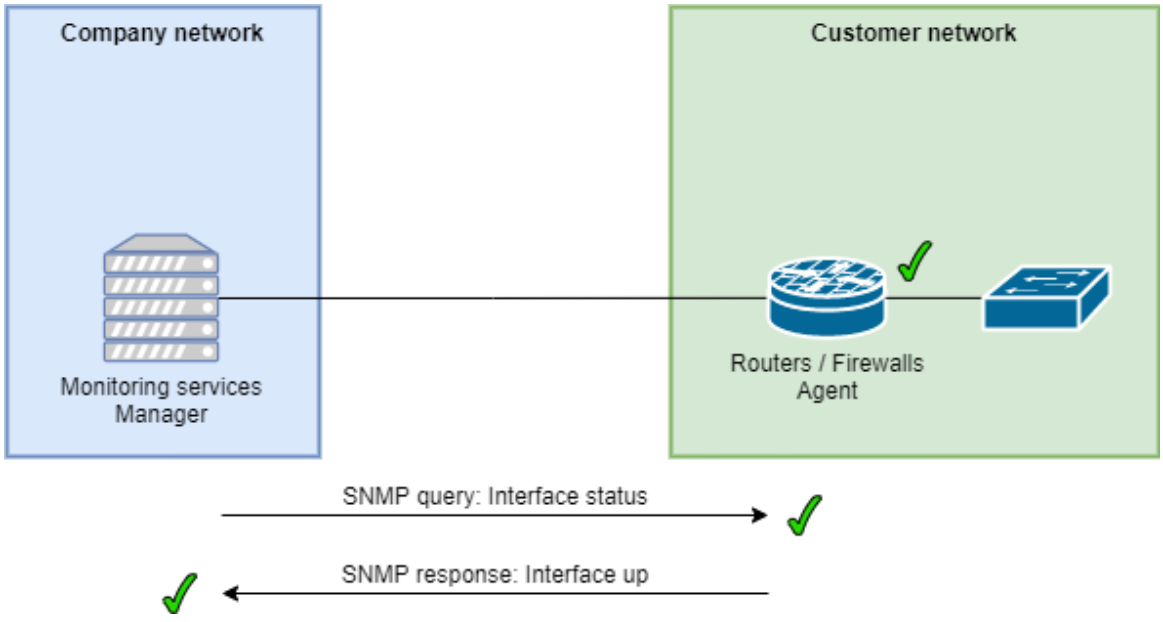

Figure 6 - SNMP query from SNMP manager for device interface status

SNMP utilizes manufacturer MIBs where components of the device are identified with Object Identifiers. Each OID is unique as stated in RFC3061 (2001, 3) and RFC1157 (1990, 12). The information for each component is stored on the OID variable which is processed by the SNMP manager.

The SNMP managers are configured with severity for each OID which determines actions for the received information as shown in figure 7. OIDs configured with the highest priority are raised to monitoring and ticketed, while the lowest priority events are discarded. Events with other priorities are added to the customer network activity reports which may be used to detect problems on the network.

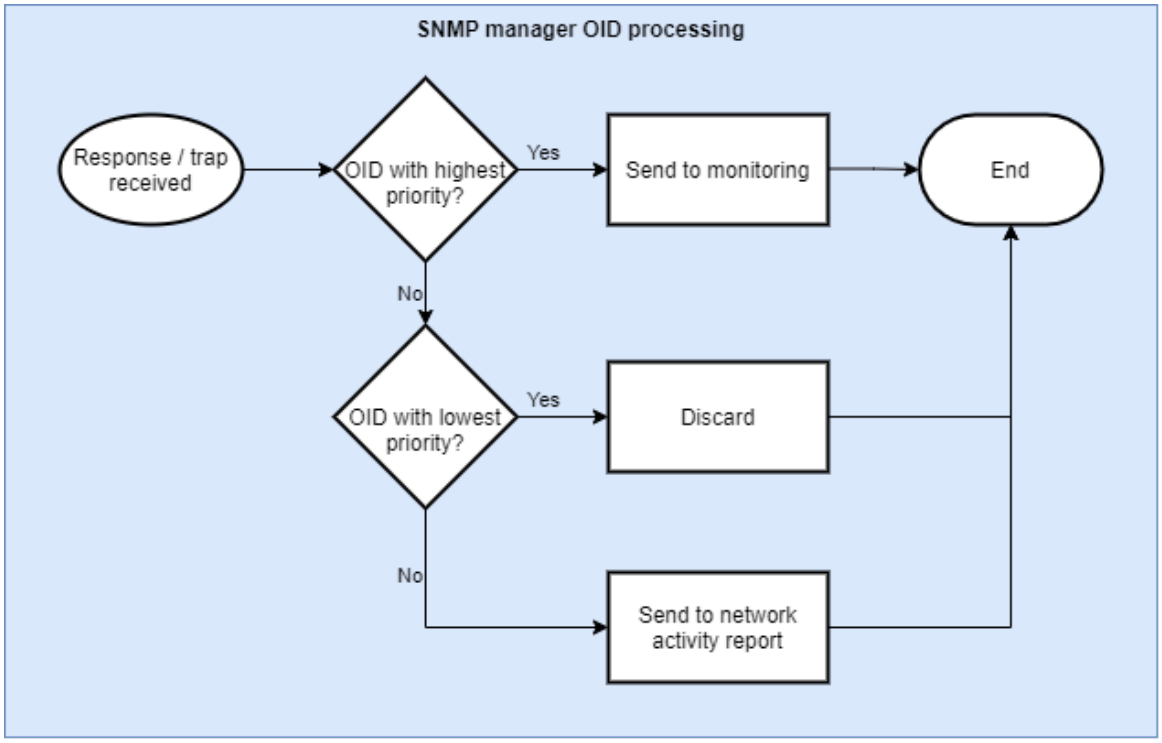

Figure 7 - SNMP manager OID priority processing

The SNMP agent may also be configured to send SNMP traps to the manager. The manager will process the received information with the same process as in figure 7 with SNMP responses. Only the highest priority traps will be forwarded to monitoring.

Figure 8 shows the router SNMP agent sending an SNMP trap from downed interface. If this port is configured to be monitored, the OID priority will be high on the SNMP manager and the event is raised to monitoring. The event will be ticketed and further processed by an employee.

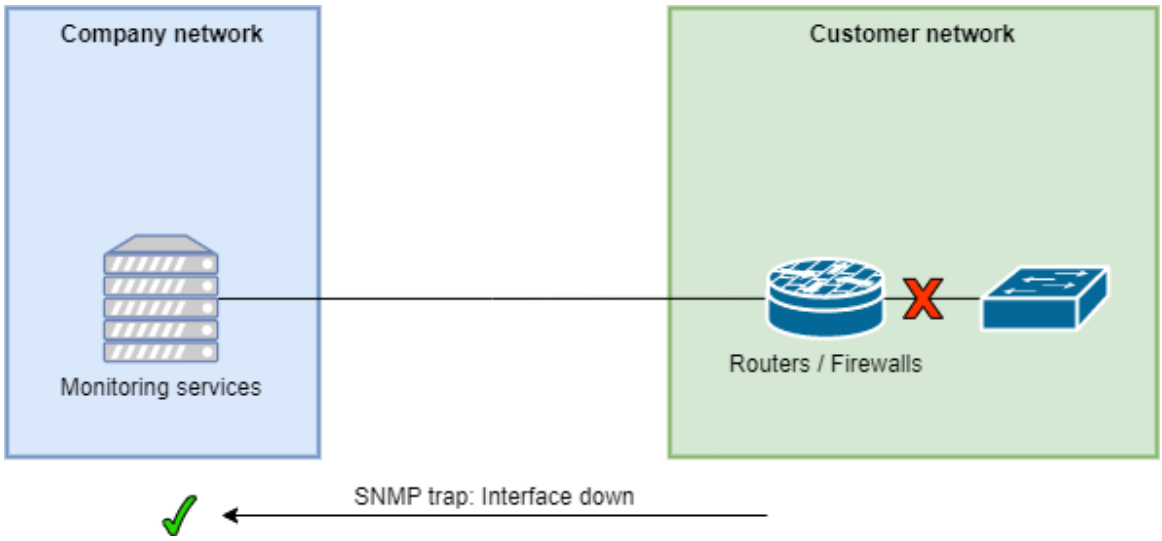

Figure 8 - SNMP trap sent to SNMP host

SNMP may also be used to create hardware monitoring if the monitoring software supports it. The network device manufacturer may include CPU or port usage OIDs in their MIBs, which can be requested to receive current usage values from the device. The monitoring software will then raise an event once usage value goes over configured threshold. These values may also be used to generate usage graphs and could help determine possible device failures or inform if the device is capable of handling traffic in its current position.

3.3 Syslog

Syslog is defined with RFC3164 but is updated with RFC5424 (2009). By default syslog messages are sent using UDP port 514 (RFC5426 2001, 5), but it is also possible to use TCP with TLS (RFC5425 2009) if the device supports it. Syslog architecture consist of an originator and collector (RFC5424 2009, 5).

The originator refers to the device generating the syslog messages and the collector to a device which receives them. The collector in spoken terms is

considered to be a syslog server which stores all the syslog messages from network devices. Figure 9 shows an example where the router sends a syslog message to the syslog server from a downed interface.

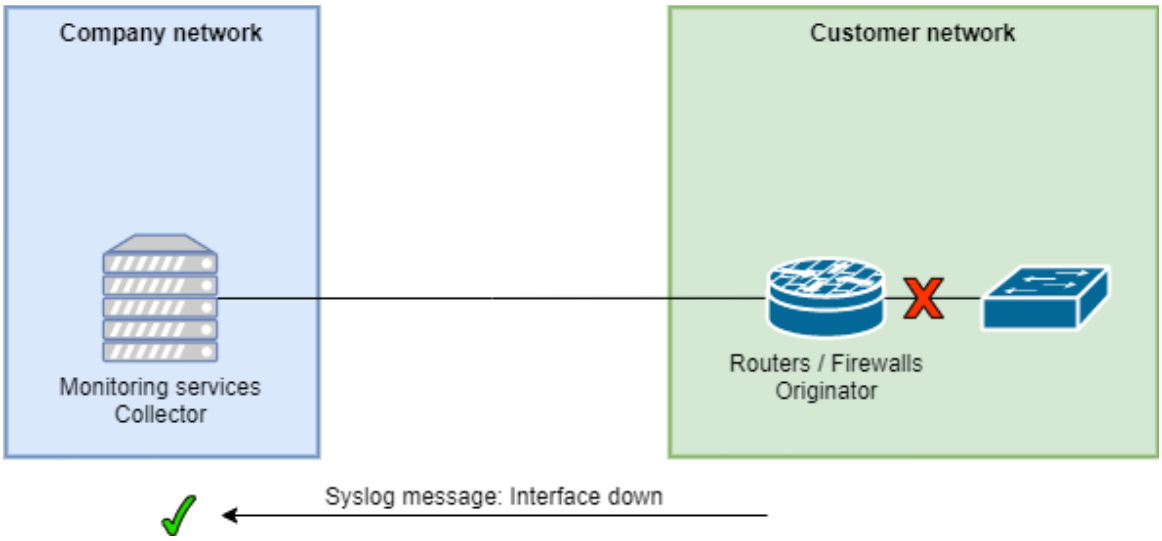

Figure 9 - Syslog message sent from an originator to collector

As this communication is done with connectionless UDP, it is not verified whether the server has received the message. Once a message is lost in transit it will not be sent again. Therefore syslog monitoring should not be considered an effective way of network monitoring, but it can be used as an additional tool with other monitoring solutions.

Syslog monitoring analyses syslog messages on the syslog server. Depending on the server software, the server can be configured to raise an event if certain keywords are detected. This method can be used as workaround for SNMP monitoring if the manufacturer MIB does not include required event for the device, or if the device status is required to be investigated after a certain syslog message.

4 PRODUCTION ENVIRONMENTS

After the company created automated production environment, the customers were being transferred to it. The transfers were done in quarters to allow production to adjust to changes with the new environment. There were mainly problems with resource management, as employees had to be trained and manage two production environments at once.

One shift consists of 3–15 employees with the shift lasting for 8–12 hours. Each of these employees are responsible for receiving calls and process incident tickets, and occasionally handle customer service requests. Only if the employee is currently working on a high-priority ticket or working with a business-critical role they will not be doing the same tasks as the rest of the shift.

4.1 Old production

The old production monitoring consists of manual work. One employee tracks events on a monitoring screen and creates tickets according to event and device priority. Each event is required to be verified by the monitoring employee before ticketing. With large network outages the monitoring screen could flood with events, which requires additional resources to process all the events within an agreed event SLA.

The monitoring employee is strictly limited to monitoring during the day and cannot handle tickets or answer calls at the same time. This role also needs to always be filled, which limits resources in the shift to actively process tickets. All generated tickets are required to be processed by an employee, which could lead tickets with cleared events remain unprocessed creating unnecessary SLA breaches.

Figures 10 and 11 show examples of event and ticket processing. The monitoring employee verifies the events and creates tickets from active events within a customer event SLA. The rest of the shift processes the tickets according to customer resolution SLA.

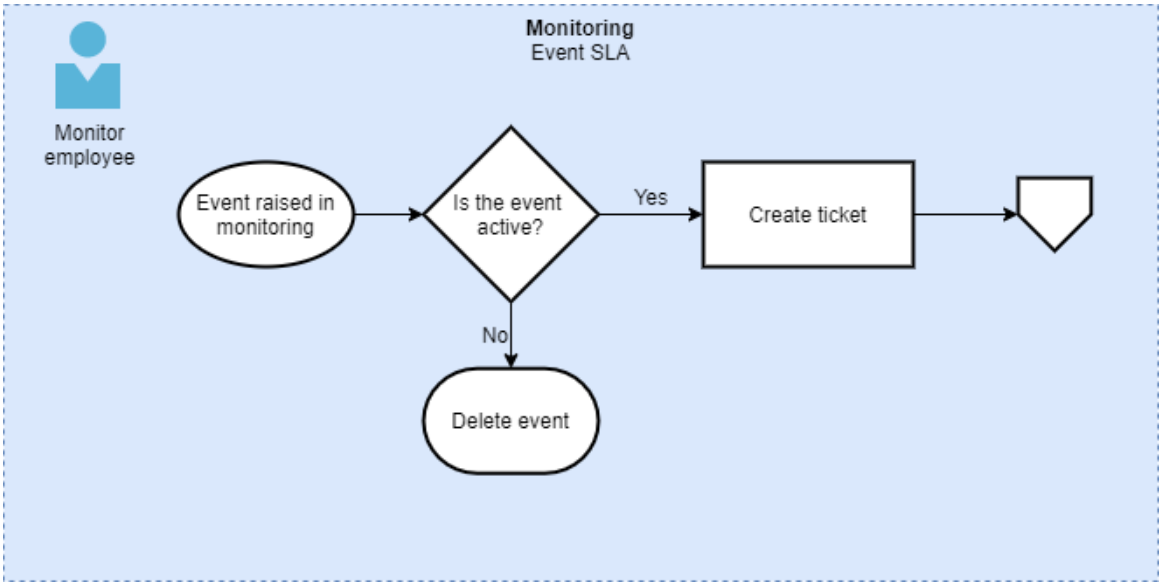

Figure 10 - Event processing

Ticket processing begins from verifying the event status and then troubleshooting the issue remotely. If the event cannot be resolved, then the customer or third party is contacted. These could vary from contacting the local contact at the customer site to verify the device status or contacting the manufacturer for technical assistance or an RMA case.

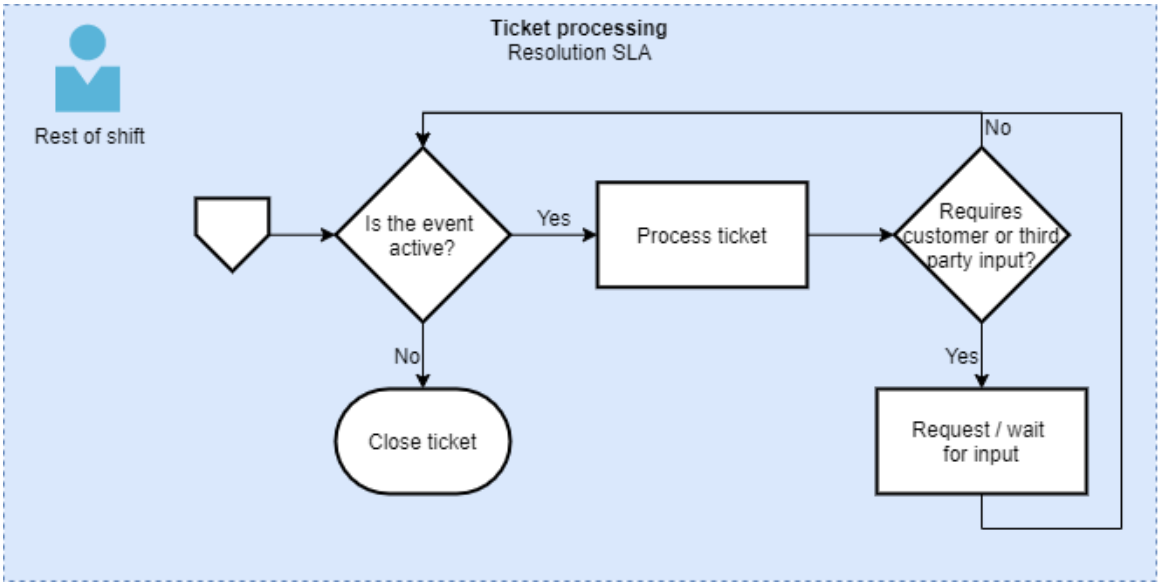

Figure 11 - Ticket processing

In old production customer documentation may be separated. Employees had to access customer documentation hosted on the company or customer servers. Updates to the documentation may require manual uploading to several locations so that each party will operate with the latest information.

Tickets also did not have clear SLA times visible in the ticket view. If the resolution SLA had to be verified, this needed to be checked from customer documentation. Also if employees required to check possible fixes from earlier tickets, they had to search for these by knowing parts of the ticket subject line or knowing the ticket ID entirely.

However this environment was familiar for employees. When working with tickets the tools were simple, and the tickets were grouped for each customer under a single customer ID. If the ticked had to be switched from incident management to some other management process, the employees could directly create a new ticket to a different management group. These operations were intuitive for the employees to use, which helps with ticket processing.

4.2 New production environment

With new production employees focus on handling tickets. The monitoring for events is automated, and tickets with cleared events closed automatically. The ticket view had remaining SLA counter visible which helps to process tickets within the agreed resolution SLA times. The search functions were improved, which made finding possible solutions faster for the detected event. Customer documentation was also made directly accessible from the ticketing portal without needing to access external sites.

Otherwise the ticket processing remains the same as in figure 11. The employee verifies the event and if it cannot be resolved remotely, the customer or third party is being contacted. Once the detected event is resolved the ticket can be closed. The customers also received portal access to the tickets, which allows them to leave comments, open new tickets or mark them as unresolved if there were problems detected at the site.

However the employees could not open tickets as freely anymore for different management process. For certain ticket types it was required to request the customer to open the ticket with so that the event processing could continue. Customers were also divided under multiple IDs, which caused confusion and required the employees to verify that the tickets were opened for the correct customer ID. On the other hand, this allows the company to create custom parameters for each ID which may help to provide personalized services for customers in their different departments.

5 BENEFITS OF AUTOMATION

The automation was to benefit production by taking care of the network monitoring and closing tickets automatically with cleared events. However at the time of writing reliable comparison could not be done. The company had transferred one of the customers to be monitored by automation, but the monitoring still required adjustment until it was working properly.

This generated large amounts of false events which were ticketed, but the tickets were also automatically closed once the monitoring was fixed. This creates a false view of the meaningful events that are detected in the customer environment. Getting an understanding of the benefits for ticket processing would require more time to gather meaningful data to understand the impact how the automation helps the production.

However, the current ticket amounts for the customer were still gathered from both production environments. It was still possible to see how many tickets were being processed by employees. First, overall ticket amounts were gathered from the year 2020 to get an understanding of average tickets per month. For comparison between production environments tickets were gathered from 2021 January for old production and 2021 April for new production. These were selected as January was the last full month with the customer being monitored in the old production, and April for the new production as it was the first full month for the customer to be monitored by automation.

The company received over 34,000 incidents during the year 2020 in old production. From these 2,700 tickets were from the customer, which totals 8% from the total ticket amount as shown in figure 12. As these were in the old production all tickets had been processed by an employee.

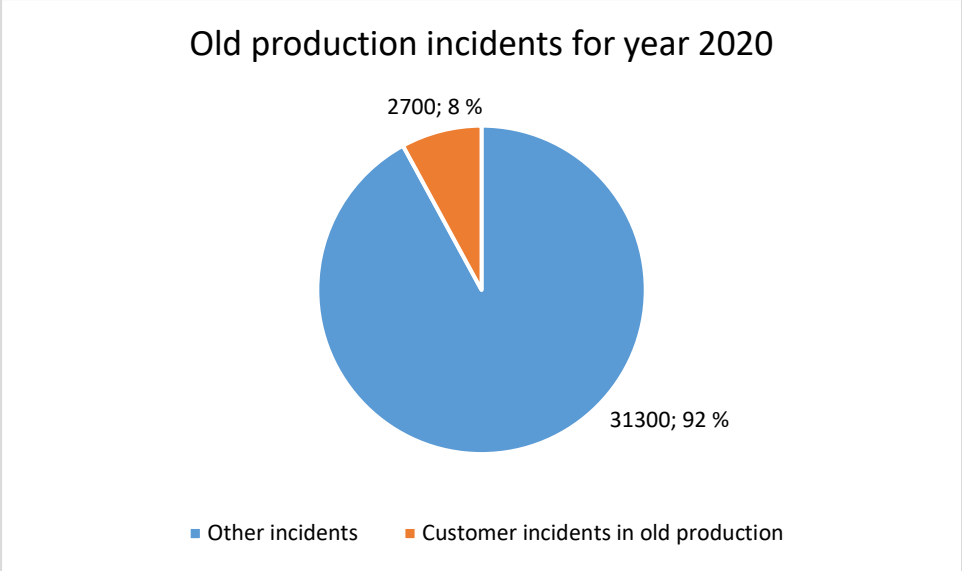

Figure 12 - Total incident tickets for the year 2020 in old production

This leads to average 225 incidents per month for 2020. It should be noted that in practice these monthly values are nowhere near static, as device failures may be unpredictable. The customer actions may also have increasing effect on the ticket amounts, as there could be maintenance work or changes with customer site infrastructure.

Figure 13 shows the opened tickets in both environments. In January the customer had 280 incidents opened at old production. New production had 840 incidents opened in April. The automation had created three times more tickets, which is expected to be caused by the errors in monitoring configuration.

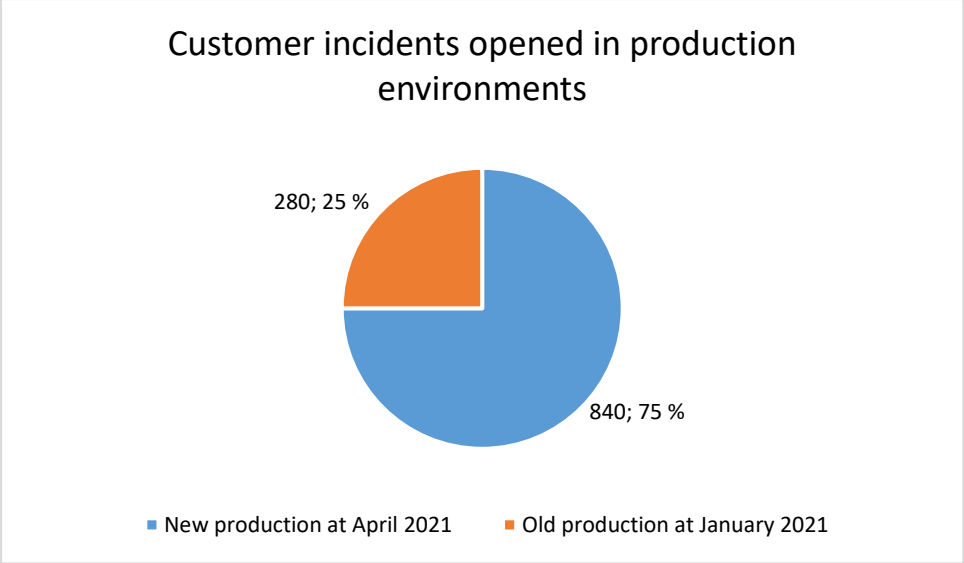

Figure 13 - Incidents opened in both production environments

From tickets in new production 195 were required to be processed by an employee, and the remaining 645 tickets were closed by automation. This totals for 23% tickets to be processed by an employee as shown in figure 14, but the percentage is suspected to be higher if the monitoring would have been configured properly. When comparing the amount of employee processed tickets in new production to the ticket average last year, this is only 30 tickets less than the counted average.

When the same comparison is done with old production in January 2021, there is a difference of 85 tickets. But comparing single months is prone to fluctuation with ticket amounts, which could lead to entirely different results next month. It would be advised to use an average when possible to create better comparisons, therefore the comparison should be done after there is enough data to find possible average to compare with.

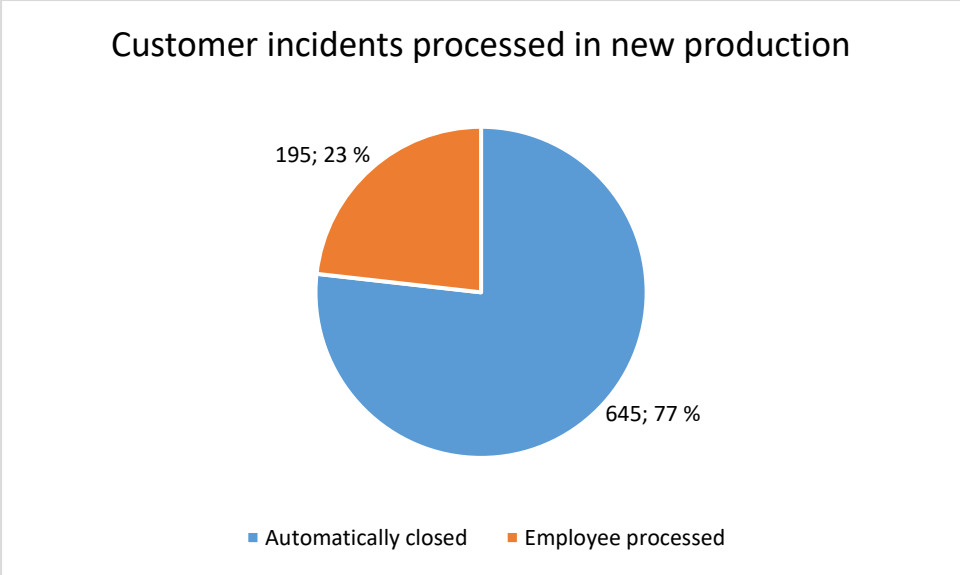

Figure 14 - Tickets processed in new production

6 CONCLUSIONS

From the comparison statistics it is early to say how well the automation will benefit the company, as valid comparison could not yet be made. With this short timespan the comparison does not give meaningful results, as the automation had created large amounts of tickets from false positive events which will affect the calculated percentages. Longer timespan would aid to reduce the noise in ticket data, as measuring with the average will not be as susceptible to fluctuation as using the ticket amounts from a single month. Additional metrics could also be included in comparison, such as ticket SLA breaches and feedback surveys from customers and employees.

The benefit of freeing the monitoring employee cannot be done until all of the customers are transferred to new production. However, the monitoring employee does not have to monitor transferred customer events anymore, which slightly eases the workload. This could prevent the company from needing to assign additional resources for event management, which will allow other employees on the shift to focus with incident management.

Shifting production environments did pose risks for the company, as the employees had to be trained for the new production and available resources had to be split between both environments. The requirements of tracking incidents on two environments at once may leave bursts of incidents unnoticed until the very last moment. This requires the monitoring company to be swift with resourcing, so that the tickets can be processed within agreed SLA. The employees also need to become familiar with the changed processes in the new environment until ticket processing can reach its full effect.

However, the company would have required to make this shift in the future, as the old software was nearing its end of life, and manual work is not viable to provide monitoring for large customer base. The new software also comes with more features, which allows the company developers to create new solutions for monitoring and other business tasks. These could include developing automation

to make first troubleshooting steps on the networking devices after events are detected, which could solve some incidents right after the event detection.

The thesis itself failed to create proper comparison between the production environments. Other benefits could be found from the shift towards automated monitoring and new software, but these benefits were not measured on the thesis. For future studies it would be recommended to plan better metrics to be measured from both environments, and to give more time for the environment to create data which would give meaningful comparison.

REFERENCES

Haimi P., Huovinen J. 2018. Industrialized service integration: Practical Guide for Management. Espoo: Sofigate Oy.

RFC1157. 1990. IETF. Web document. Available at: https://tools.ietf.org/html/rfc1157 [Accessed 24 March 2021].

RFC3061. 2001. IETF. Web document. Available at: https://tools.ietf.org/html/rfc3061 [Accessed 24 March 2021].

RFC3164. 2009. IETF. Web document. Available at: https://tools.ietf.org/html/rfc3164 [Accessed 30 April 2021].

RFC3411. 2002. IETF. Web document. Available at: https://tools.ietf.org/html/rfc3411 [Accessed 24 March 2021].

RFC5424. 2009. IETF. Web document. Available at: https://tools.ietf.org/html/rfc5424 [Accessed 30 April 2021].

RFC5425. 2009. IETF. Web document. Available at: https://tools.ietf.org/html/rfc5425 [Accessed 30 April 2021].

RFC5426. 2009. IETF. Web document. Available at: https://tools.ietf.org/html/rfc5426 [Accessed 30 April 2021].

RFC792. 1981. IETF. Web document. Available at: https://tools.ietf.org/html/rfc792 [Accessed 24 March 2021].

Service Design. 2007. Office of Government Commerce. London: The Stationery Office.

Service Operation. 2007. Office of Government Commerce. London: The Stationery Office.

The Official Introduction to the ITIL Service Lifecycle. 2007. Office of Government Commerce. London: The Stationery Office.