BPI Challenge 2013 - Applied process mining techniques for incident and problem management

Peter Van den Spiegel, Leen Dieltjens and Liese Blevi
KPMG Advisory, IT Advisory, Bourgetlaan 40, 1130 Brussels, Belgium
[pvandenspiegel, ldieltjens, lblevi]@kpmg.com

Abstract. The incident and problem management process forms an essential part in every organization. Since businesses rely heavily on IT, each outage, issue or user service request should be dealt with as quickly as possible in order to minimize its impact on operations. For Volvo IT Belgium, we analyzed the event log file of an incident and problem management system called VINST, in order to objectively verify the efficiency and effectiveness of the underlying process. Our analysis was performed by means of a of a combination of process mining and data mining techniques and tools, including Disco, ProM, Minitab and MS Excel. The log file itself consisted of 65.533 incident records and 6.660 problem records. As part of the exercise, we investigated aspects, such as total resolution times of tickets, actual resolution process being followed, ping-pong behavior between the different helpdesk lines, differences between distinct support teams etc. Finally, we also made recommendations to improve the current process and increase integration between incident and problem management.

Keywords: BPI Challenge, Process Mining, Data mining, Incident management, Problem management

1 Introduction

Within the IT Advisory department of KPMG Belgium, we help clients bridging the GAP between business and IT and in doing so assist them to leverage their IT and information assets to maximize business value. A part of our services involves the optimization of the internal IT organization, which includes incident & problem management. Another service relates to data analytics on business processes, which we have been applying for several years, in both our audit and advisory engagements. A few months ago, we came across "process mining" and quickly discovered the additional value it could bring to the way we were analyzing and reporting data. Having heard of this BPI Challenge 2013, we believed this would give us the perfect challenge to put these techniques in practice.

2 Executive Summary

Based on our analyses performed, we identified several areas for improvement within the current incident and problem management processes, in particular with respect to the overall through-put time (time between registration and closing of incident), the wait-time and the ping-pong behavior. Although a reasonable percentage of incidents was solved in first line (i.e. 60%), we noted that the through-put time (time between registration and closing of incident) appears to be longer than 10 days, in more than 31% of the cases. 12% of all incidents have been open for more than 20 days.

As the incident process is supposed to be resolving incidents as quickly as possible (e.g. via quick-fix or work-around), we believe that the way in which the incident process is currently being executed, could have a negative on the business operations. We are of the opinion that through better assigning the correct support team and thus limiting the number of ping-pong taking place, limiting or closely monitoring the use of status "waiting.." as well as aligning IT departments A2 and C could significantly improve the overall incident management process and business support.

We also recommend linking the incident and problem management process, in order to evaluate to what extent recurring or critical incidents are properly handled within the problem management process. We noted 819 open problems. However, as we did not have sufficient data available to investigate whether a) incidents are recurring and b)incidents are "resulting" into problems, we were unable to evaluate the effectiveness of the problem management process.

Below, we provide an answer on the key questions as asked by the process owner:

- **Push to Front**

We noted that the majority (60%) of the incidents have been solved by the first line. However, the IT department C had a significant higher push-to-front (68,85%) compared to only 22,8% for organization A2. We also noted that 17% of the incidents managed in second line, have not been initiated by a first line, but were immediately handled in second line. The product which was most appearing in the incidents, both in first line as in second line is Product 424 with respectively 15% of all incidents in first line and 7,82% of all incidents in second line. Looking at the problem list, we noticed that only 1,41% of the problems was related to this product. This leads us to conclude that incidents regarding this product might not be sufficiently picked up within the problem management process (in order to find a permanent fix).

- **Ping Pong Behavior**

Our analyses showed significant evidence of ping pong behavior amongst teams. For example, we noted that for calls solved within the first line, only 72% have been solved by the initially appointed team. Overall this percentage was 49%. In 23,5% of the cases the incident was passed on to another team. For 27%, two or more movements were involved. On average we noted that 2,25 support teams are involved per incident.

- **Wait User abuse**

Based on our analysis, More than 34% of the total through-put time is caused by “wait-user”. The “wait-user” is particularly used in the first line and lasts more than 1 week in more than 29% of the cases. This gives us indication that the “wait…” status might be abused to reduce the actual resolution time. Waiting time in relation to the total through-put time is higher for C (38,45%) compared to A2 (28,35%).

- **Process Conformity per Organization**

Clear differences were noted between the way both IT organizations execute their incident management processes. Within organization C, the percentage push-to-front is significantly higher (i.e. 68%) than within organization A (22,8%). We also noted that although the through-put time for both Organisation A2 and C is similar (median is respectively 8,27 days and 7,48 days), the variation of through-put time is significantly higher for A2 (standard deviation 57,46 days) compared to C (standard deviation of 22,52 days). Based on our analyses we noted that overall predictability of the incident handling process for organization C is higher than it is for organization A2. Moreover, given the significant higher push-to-front for C compared to A2, we are of the opinion that the organization C is performing better than organization A2 (with the exception of the use of waiting time).

3 Understanding the process

3.1 Mapping statuses on the standard process flow

On the basis of the general process information provided by Volvo IT Belgium and the Information Technology Infrastructure Library (ITIL v3) we determined the standard process flow (Figure 1).

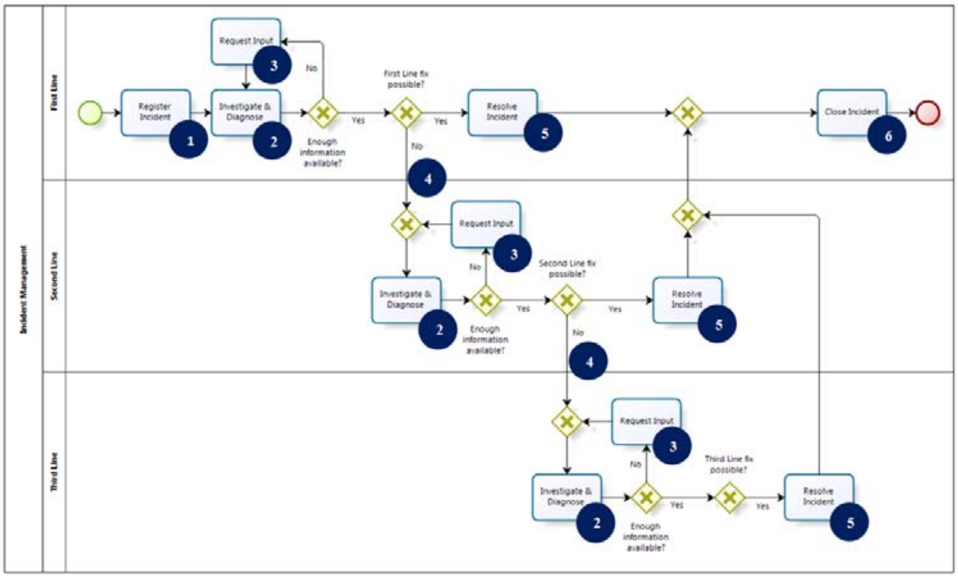

Fig. 1. Standard process flow.

In our standard process flow, we identified 5 main activities:

- **Register incident:** logging, categorization and prioritization of the incident
- **Investigate & Diagnose:** research of the issue to determine cause and remediation options
- **Request input:** request for input from customer, user, vendor, etc. in order to continue investigation and diagnosis
- **Resolve incident:** implementation of the solution or workaround in order to restore service
- **Close incident:** verification whether service has restored and closure of the incident

To facilitate the understanding of the process and further analysis, we mapped the 13 statuses on the main activities. The mapping is based on the description of the statuses, as listed in Table 1. We assumed that the status of an incident is changed upon execution of an activity. For example, if an activity owner completes activity *Register Incident* the status of the incident is changed to Accepted/Assigned. We marked the moments of status change on the process flow in Fig. 1. The same logic is applied for the other statuses.

Table 1. Mapping of the statuses on the standard activities.

Activity	Status Icon	Status	Description
Register Incident	1	Accepted/Assigned	The incident is assigned and acknowledged by the suggested Support Team (ST).
Investigate & Diagnose	2	Accepted/In Progress	The incident is acknowledged and currently being worked on by the ST.
Request Input	3	Accepted/Wait	The incident is acknowledged, but input is requested from a third party in order to diagnose the issue.
		Accepted/Wait-User	The incident is acknowledged, but input is requested from the User in order to diagnose the issue.
		Accepted/Wait-Customer	The incident is acknowledged, but input is requested from the Customer in order to diagnose the issue.
		Accepted/Wait-Vendor	The incident is acknowledged, but input is requested from the Vendor in order to diagnose the issue.
		Accepted/Wait-Implementation	The incident is acknowledged, but cannot be solved immediately because of implementation restrictions.
Resolve incident	5	Completed/In Call	A solution is found and implemented during call.
		Completed/Resolved	A solution is implemented.
		Completed/Cancelled	The incident is cancelled. No solution needs to be implemented, so the incident can be considered as resolved.
Close incident	6	Completed/Closed	The solution is verified and the incident is closed.
Other	4	Queued/Awaiting Assignment	The incident cannot be solved by the assigned ST and is transferred to another one.
		Unmatched/Unmatched	The incident could not be matched to existing incidents in the system.

Remark 1. An incident can be reassigned multiple times, also within the same support line. This means that the status of an incident can be changed to Queued/Awaiting Assignment anywhere in the process flow.

3.2 Main scenarios

From the standard process flow, we derived 5 main scenarios:

1. Incident is solved in First Line (normal flow)
Register Incident → Investigate and Diagnose (→ Request input) → Resolve Incident → Close Incident

2. Incident is not closed yet
Register Incident → Investigate and Diagnose (→ Request input) → Resolve Incident

3. Incident is reopened
Register Incident → Investigate and Diagnose (→ Request input) → Resolve Incident → Close Incident → Investigate and Diagnose

4. Incident is transferred to Second Line
Register Incident → Investigate and Diagnose (→ Request input) → Queued/Awaiting Assignment → Investigate and Diagnose → Resolve Incident → Close Incident

5. Incident is transferred to Third Line
Register Incident → Investigate and Diagnose (→ Request input) → Queued/Awaiting Assignment (x2) → Investigate and Diagnose → Resolve Incident → Close Incident

In our analysis we will not focus on scenario 2 and 3. We decided to focus on the questions that were asked by the process owner, which can be linked to scenarios 1, 4 and 5.

3.3 Link between incidents and problems

According to leading practices in incident and problem management processes, there should be a close relationship between incidents and problems. Whereas incident management primarily focuses on helping the end-user as quickly as possible, the problem management process investigates the root cause of recurring incidents to provide a long term solution. Known root causes and their solutions should be entered in a known-error database (KeDB), allowing the service desk to reach a higher percentage of first-line fix and higher resolution times.

4 Pre-assessment of the data set

Before analyzing the process, we performed a pre-assessment of the data set, in order to increase our understanding of the data received.

We generated a frequency table based on the creation dates of the tickets (i.e. the moment a first registration in the system is available) until the closure of a ticket (i.e. status completed)

Table 2. Frequency table (all tickets)

Period		Total Created	Total Closed
…	1/01/2010	**0**	**0**
2/01/2010	1/07/2010	**1**	**0**
30/06/2010	1/01/2011	**0**	**0**
31/12/2010	1/07/2011	**17**	**0**
30/06/2011	1/01/2012	**44**	**0**
31/12/2011	1/07/2012	**7491**	**7553**
30/06/2012	1/01/2013	**0**	**0**
31/12/2012	1/07/2013	**0**	**0**

Table 3. Frequency table (tickets 2012)

Period		Total Created	Total Closed
…	1/01/2012	**62**	**0**
2/01/2012	1/02/2012	**43**	**0**
2/02/2012	1/03/2012	**45**	**0**
2/03/2012	1/04/2012	**176**	**0**
2/04/2012	1/05/2012	**2029**	**19**
2/05/2012	1/06/2012	**5198**	**7533**
2/06/2012	1/07/2012	**0**	**1**
2/07/2012	1/08/2012	**0**	**0**
2/08/2012	1/09/2012	**0**	**0**
2/09/2012	1/10/2012	**0**	**0**

Given the apparent low number of tickets created in the preceding months, compared to the tickets created in May, we believe that tickets that had been created and closed before May 2012, have not been taken into account for the download. As a result, we believe there is a risk that our analysis might not represent the actual situation in terms of resolution times, involved support teams and comparison of performance between teams.

5 Analysis of the process

In our analysis of the process we focused on the questions formulated by the process owner.

1 **Push to Front (incidents only)** Is there evidence that cases are pushed to the 2nd and 3rd line too often or too soon?
2 **Ping Pong Behavior** How often do cases ping pong between teams and which teams are more or less involved in ping-pong?
3 **Wait User abuse (incidents only)** Is the “wait user” substatus abused to hide problems with the total resolution time?
4 **Process Conformity per Organisation** Where do the two IT organisations differ and why

We used a combination of different tools to perform our analysis. We used the demo version of Disco that was provided by Fluxicon, the open source process mining tool ProM 5.2, Microsoft Excel and MiniTab.

5.1 Push to Front

The main objective of incident management is helping the end-user as quickly as possible. To reach this objective it is important to have a good push to front process. We evaluated the push to front process by answering the following questions:

1. *How many incidents are resolved in First Line, without interference of a Second or Third Line Support Team?*
2. *Where in the organization is the push to front process most implemented, specifically if we compare the Org line A2 with the Org line C?*
3. *For what products is the push to front mechanism most used and where not?*
4. *What functions are most in line with the push to front process?*

5.1.1 How many incidents are resolved in First Line, without interference of a Second or Third Line Support Team?

To answer this question, we used Disco's built-in filtering algorithm. First of all we removed all open incidents by setting an Endpoint filter on Activity. The result of this filter was a reduction of less than 1% of the total number of cases. In absolute figures we noted that out of the 7.554 incidents, 7.546 are completed. By setting an additional Attribute filter on org:group (Involved ST), we were able to exclude all cases that were completed with interference of a Second and/or Third Line Support Team. More specifically, as shown in Figure 1, we used filtering mode Forbidden to remove all cases that have an org:group with 2nd or 3nd in their name.

By applying the filter algorithm, we identified that 60% of the incidents was closed in First Line. The variant statistics in Disco show us the following:

- Among the 4.542 incidents closed in First Line there is a total of 942 process flow variants.
- The most common process flow variant, representing 37,87% of the cases, goes through the following 3 statuses in sequence: Accepted/In Progress, Accepted/In Progress and Completed/In Call.
- The most common process flow variant has a mean duration of 38 minutes and 4 seconds.

We noted that 72% of the incidents is resolved by the initially assigned support team.

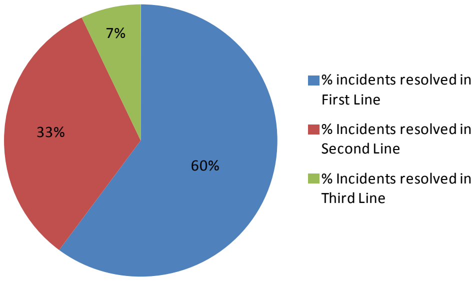

Fig. 2 Distribution across support lines

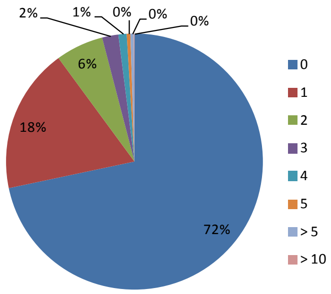

Fig. 3 Number of transfers between ST's

We compared the mean throughput time of the incidents across the different support lines. As shown in the Fig. 4 below, the throughput time increases significantly once the incident is transferred to Second Line.

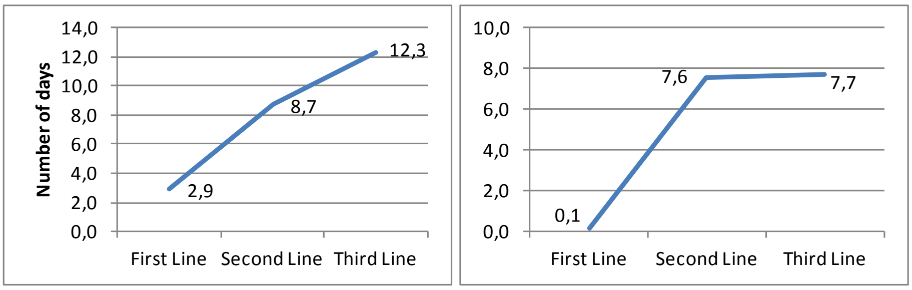

Fig. 4 Throughput time per support line *(right: total population, left: only may)*

5.1.2 Where in the organization is the push to front process most implemented, specifically if we compare Org line A2 with Org line C?

Table 4. Comparison of the First Line Level Fix between Org line A2 and Org line C

Metrics	Org line A2	Org line C
Number of incidents resolved	744	5746
Number of incidents resolved in First Line	171	3956
Percentage of incidents resolved in First Line	22,98%	68,85%

If we compare the metrics, we can conclude that Org line C resolves the most incidents without interference of Second and/or Third Line. Therefore we can say that Org line C is most in line with the Push to Front process.

5.1.3 For what product is the push to front mechanism most used and where not?

Table 5. Number of products affected across the different support lines

Metrics	

Total number of products	705
Total number of products handled in First Line	368
Total number of products handled in Second Line	474
Total number of products handled in Third Line	162

Table 6. Overview of the most affected product per support line

Support Line	Most affected product	Absolute occurence	Relative occurrence within the support line	Relative occurrence within the whole support org.
First Line	PROD424	684	15,06%	9,06%
Second Line	PROD424	193	7,82%	2,56%
Third Line	PROD607	69	12,85%	0,91%

In Second Line we see the highest variety in affected products. We noted that in both First Line and Second Line most incidents are related to product PROD424. In Third Line, product PROD607 is most affected.

If we compare the products across the different support lines, we note that incidents related to product PROD566 are always resolved in First Line. In the tables below an overview is given of the products that are most affected in First Line and an overview of those products that are always solved in First Line.

Table 7. Overview of the most affected products in First Line (Left: all products, Right: products that only handled in First Line)

Product	Absolute occurrence	Relative occurrence	Product	Absolute occurrence	Relative occurrence
PROD424	**684**	**15,06%**	PROD566	**158**	**3,48%**
PROD660	**442**	**9,73%**	PROD328	**40**	**0,88%**
PROD383	**193**	**4,25%**	PROD832	**39**	**0,86%**
PROD253	**172**	**3,79%**	PROD369	**30**	**0,66%**
PROD566	**158**	**3,48%**	PROD505	**20**	**0,44%**
PROD494	**142**	**3,13%**	PROD420	**19**	**0,42%**
PROD13	**107**	**2,36%**	PROD522	**15**	**0,33%**
PROD321	**94**	**2,07%**	PROD732	**15**	**0,33%**
PROD267	**79**	**1,74%**	PROD533	**14**	**0,31%**
PROD453	**77**	**1,70%**	PROD794	**14**	**0,31%**

5.1.4 What functions are most in line with the push to front process?

Table 8. Overview of functions involved in First Line

Function	Absolute

	occurrence
V3_2	3533
A2_1	498
E_5	256
A2_5	11
E_6	2
A2_2	1
E_10	1

We identified that across the 20 functions, 8 are involved in incidents resolved in First Line. Based on the figures as shown in the table above, function V3_2 is most in line with the push to front process. We noted that out of the 4804 incidents handled by V3_2, 3533 are resolved in First Line.

5.2 Ping Pong Behavior

Questions:

1.1 *What are the support teams that are responsible for most of the ping pong?*
1.2 *What are the functions that are responsible for most of the ping pong?*
1.3 *What are the organizations that are responsible for most of the ping pong?*
1.4 *What products are most affected by it?*

5.2.1 INCIDENTS

5.2.1.1 What are the support teams that are responsible for most of the ping pong?

Out of 649 different support teams, we noticed that three teams stand out the most in effectuating the most statuses: G97, G96, S42. When taking a look at the top ten, we noticed that almost all of them are part of the first line support. The average number of statuses where each support team is involved in, is about 100.

Table 1. Support team metrics

Support team	*Relative occurrence (%)*
G97	11,39
G96	9,15
S42	6,68
G230 2nd	2,53
D5	2,53
S56	2,42

D8	2,41
G76	1,98
D2	1,86
G92	1,76

We considered these ten support teams in our analysis in ProM, using the social network miner, and we saw that some of these support teams hand over the work to each other.

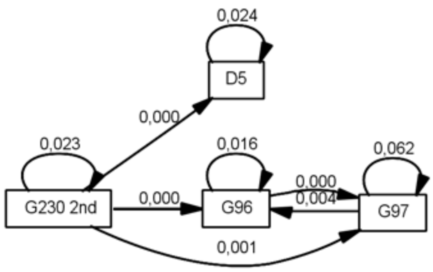

Fig. 5. Handover of work.

G96 and G97 more often exchange the work, in comparison to the others. And G96 is very popular for receiving work from others.

When we applied this mining technique on all support teams, we discovered the following regarding the handover of work: (threshold = 0)

Fig. 6. Handover of work

G96 and G97 are the most popular support teams in receiving work from other teams. The only circular relationship that was found is the one between G96 and G97.

5.2.1.2 What are the functions that are responsible for most of the ping pong?

In considering the function division of the support team we noticed that this is not always filled out. We filtered out these blank function division lines and only selected those service requests where G96 and G97 are involved in. Obviously, function division V3_2 is responsible for most of the Ping Pong. Also note that multiple support team function divisions are active within one service request.

Table 9. Support team function metrics

Support team function division	*Relative occurrence (%)*
V3_2	99,52
E_10	0,04
A2_4	0,22
A2_2	0,04
A2_1	0,17

5.2.1.3 What are the organizations that are responsible for most of the ping pong?

When selecting only those service requests with G96 and/or G97 involvement, we can state that Organization Line C is affected the most by the ping pong behavior.

Table 20. Involved Organisation Line metrics

Involved Organization Line	*Relative occurrence (%)*
Org line C	80,38
Org line A2	9,06
Org line B	8,62

5.2.1.4 What products are most affected by it?

For the products that are affected the most with the Ping Pong behavior, we selected only those service requests with G96 and/or G97 involvement;

Table 31. Support team function metrics

Product	*Relative occurrence (%)*
PROD424	27,87
PROD494	4,5
PROD698	4,2

Product ‘PROD424’ is obviously the most affected.

5.2.2 OPEN PROBLEMS

5.2.2.1 What are the support teams that are responsible for most of the ping pong?

We noted that there are 187 different support teams that contribute to the open problems. The three teams that issue the most statuses are G42 3rd, S33 2nd and G88 2nd. In the top ten of this list are also no first line support teams. The average number of statuses where each support team is involved in, is about 12, and the median is 4.

Table 42. Support team metrics

Support team	*Relative occurrence (%)*
G42 3rd	7,36
S33 2nd	7,19
G88 2nd	5,87
G199 3rd	5,27
G273 3rd	4,34
G271 2nd	4,25
S30 2nd	4,25
G55 2nd	2,81
M1 2nd	2,59
G230 2nd	2,51

If we use the social network miner in ProM (threshold = 0), we noticed that there is a handover of work from the third line support team to the second line support team: from G273 3rd to G88 2nd.

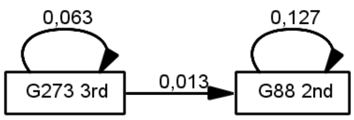

Fig. 7. Handover of work.

When we applied this mining technique on all the support teams, we discovered the following handover of work; (threshold = 0)

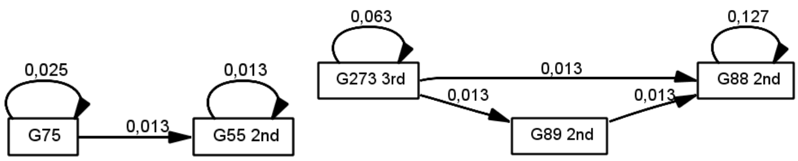

Fig. 8. Handover of work.

G88 is very popular for receiving work from other support teams. G273 3rd tends to hand over its work instead of receiving work from other teams. No circular relationship was found.

5.2.2.2 What are the functions that are responsible for most of the ping pong?

With the help of an SQL query we extracted the problem requests with involvement of the support teams that were involved in the handover of work. Support team function division E_4 is then responsible for most of the ping pong.

Table 53. Support team function metrics

Support team function division	*Relative occurrence (%)*
E_4	30,94
C_6	23,57
E_10	20,49
E_1	13,52
E_5	5,53
E_6	1,64
A2_1	0,82
C_1	0,82
E_7	0,61
E_8	0,61
A2_2	0,41
A2_3	0,41
V3_3	0,41
V3_2	0,20

5.2.2.3 What are the organizations that are responsible for most of the ping pong?

Using the output of the same SQL query we can count which Organization Lines have been responsible for most of the Ping pong behavior: Org line C.

Table 64. Involved Organisation Line metrics

Involved Organisation Line	*Relative occurrence (%)*
Org line C	97,35
Org line A2	1,63
Org line G4	0,61
Org line B	0,41

5.2.2.4 What products are most affected by it?

Based on the SQL query previously run, we discovered the products that are affected the most with the Ping Pong behavior: PROD802.

Table 75. Product metrics (top 10)

Product	*Relative occurrence (%)*
PROD802	38,90
PROD793	15,68
PROD745	13,03
PROD436	6,31
PROD348	6,11
PROD154	2,44
PROD327	2,44
PROD673	2,24
PROD597	1,83
PROD821	1,83

5.2.3 CLOSED PROBLEMS

5.2.3.1 What are the support teams that are responsible for most of the ping pong?

There are 130 different support teams that have worked on the closed problems. The three teams that were involved in the most status changes are G199 3rd, S33 2nd and G21 2nd. The third line support is involved in the majority of the closed problems. When we look at the top ten of the teams involved in closed problems, we notice that mostly second and third line support are involved.

Table 86. Support team metrics (top 10)

Support team	*Relative occurrence (%)*
G199 3rd	36,20
S33 2nd	5,10
G21 2nd	4,85
M1 2nd	3,79
S30 2nd	3,19
G357 2nd	3,06
G141 3rd	2,64

G42 3rd	2,08
G51 2nd	1,40
M3 2nd	1,32

If we use the social network miner in ProM (threshold = 0) for these 10 support teams, we noticed that there is a handover of work from the second line support team to the third line support team: from G21 2nd to G199 3rd.

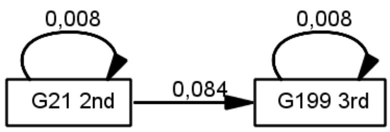

Fig.9. Handover of work.

After applying the social network miner on the whole population of support teams, we discovereed the following handover of work; (threshold = 0)

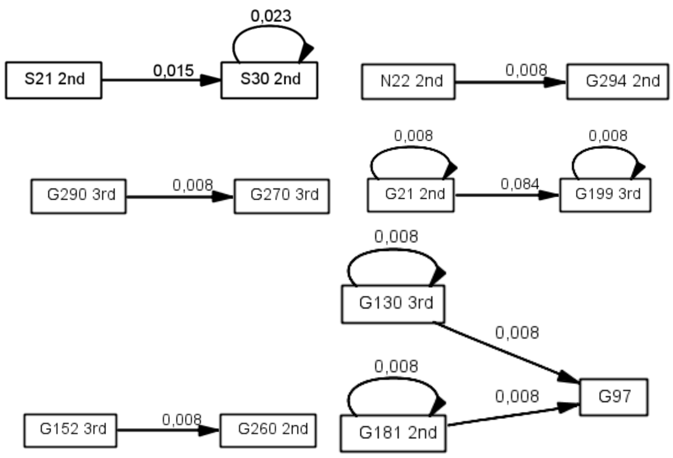

Fig.10. Handover of work.

S21 2nd and G97 are the most popular support teams to handover the work to. No circular relationship was found.

5.2.3.2 What are the functions that are responsible for most of the ping pong?

By using an SQL query we selected only those problem requests that have an involvement in the handover of work: S21 2nd, S30 2nd, N22 2nd G294 2nd, G290 3rd, G270 3rd, G21 2nd, G199 3rd, G152 3rd, G260 2nd, G130 3rd, G181 2nd, G97.

We noted that for more than 2000 status changes the support team function division was not filled out.

Support team function division C_6 is responsible for most of the ping pong.

Table 97. Support team function metrics (top 10)

Support team function division	*Relative occurrence (%)*
<BLANK>	31,20
C_6	7,99
A2_2	7,45
E_10	6,70
A2_3	6,25
E_4	6,04
A2_1	6,01
E_8	4,02
E_1	3,54
A2_4	3,06

5.2.3.3 What are the organizations that are responsible for most of the ping pong?

Using the output of the same SQL query we can count which Organization Lines have been responsible for most of the Ping pong behavior: Org line C.

Table 108. Involved Organisation Line metrics (top 10)

Involved Organisation Line	*Relative occurrence (%)*
Org line C	40,57057
Org line A2	26,51652
Org line G3	17,47748
Org line G4	9,129129
Org line B	2,612613
Org line V2	1,381381
Org line F	0,750751
Org line V11	0,495495
Org line V7n	0,405405
Org line D	0,24024

5.2.3.4 What products are most affected by it?

Based on the SQL query previously run, we discovered the products that are affected the most with the Ping Pong behavior: PROD97.

Table 119. Product metrics (top 10)

Product	*Relative occurrence (%)*
PROD97	9,19
PROD98	5,75
PROD802	4,49
PROD96	4,34
PROD374	2,54
PROD412	2,27
PROD793	2,09
PROD597	1,98
PROD660	1,83
PROD236	1,53

5.3 Wait User abuse

We understand that Volvo IT Belgium applies a lot of KPI's related to the resolution time of incidents. The use of sub status "Wait User" may have a significant impact on those KPI's. In order to provide insight in the "Wait User" usage across the organization, we provided an answer to the following questions:

1. *Who is making most use of the sub status "Wait-User" (action owner)?*
2. *What is the behavior per support team?*
3. *What is the behavior per function?*
4. *What is the behavior per organization?*
5. *Is there any (mis)-usage per location?*
6. *What is the average duration an incident is in status "Wait – User"?*

5.3.1 Who is making most use of the sub status "Wait-User"?

Table 20. General metrics

General metrics	*No of Records*
Number of records with sub status "Wait – User"	4.217
Number of owners	1.440
Number of owners who made use of the sub status "Wait – User"	580

We listed all owners with the number of times the owners used the sub status "Wait – User" versus the number of times the owners were involved in any incident. This analysis was performed to identify the "action owner – wait user ratio", i.e. what is the usage percentage of status "Wait – User" by an action owner in comparison with the total number of times the owner was involved in an incident.

We noted 143 action owners who used the sub status "Wait – User" in 20% or more of all their actions, whereby 1 action owner "Sreeraghu" used the sub status "Wait – User" in 100% of all his actions. However, this owner was only involved in 1 incident.

To eliminate this type of users, we excluded all action owners who performed less than 50 actions in order to identify those action owners who are a lot involved in incidents.

We noted 15 action owners, which were involved in 50 actions or more, who used the sub status "Wait – User" in 20% or more of all their actions. (See appendix Wait user behavior – 1)

5.3.2 What is the behavior per support team?

Table 21. General metrics

General metrics	*No of Records*
Number of records with sub status "Wait – User"	4.217
Number of support teams	649
Number of support teams who made use of the sub status "Wait – User"	298

We listed all support teams with the number of times the support teams used the sub status "Wait – User" versus the number of times the support teams were involved in any incident. This analysis was performed to identify the "support team – wait user ratio", i.e. what is the usage percentage of status "Wait – User" by a support team in comparison with the total number of times the support team was involved in an incident. Both databases were linked based on the common data field "support team".

We noted 30 support teams who used the sub status "Wait – User" in 20% or more of all their actions, whereby "G32 2^{nd}" used it the most (i.e. 9 times Wait – User sub status used with a total of 26 actions). To identify the support teams who used the sub status Wait – User significantly more, we excluded all support teams who performed less than 50 actions.

We noted 3 support teams, which were involved in 50 actions or more, who used the sub status "Wait – User" in 20% or more of all their actions: N45, L50 3^{rd} and G73. (See appendix Wait user behavior – 2)

5.3.3 What is the behavior per function?

Table 22. General metrics

General metrics	*No of Records*
Number of records with sub status “Wait – User”	4.217
Number of functions	24
Number of functions who made use of the sub status “Wait – User”	20

We listed all functions with the number of times a function used the sub status “Wait – User” versus the number of times the function was involved in any incident. This analysis was performed to identify the “function – wait user ratio”, i.e. what is the usage percentage of status “Wait – User” by a function in comparison with the total number of times the function was involved in an incident. Both databases were linked based on the common data field “function”.

We noted 11 functions who used the sub status “Wait – User” in 5% or more of all their actions, whereby “D_1” used it the most (in %: 10,82% - i.e. 161 times Wait – User sub status used with a total of 1.488 actions). (See appendix Wait user behavior – 3)

5.3.4 What is the behavior per organisation?

Table 23. General metrics

General metrics	*No of Records*
Number of records with sub status “Wait – User”	4.217
Number of organizations	25
Number of organizations who made use of the sub status “Wait – User”	15

We listed all organizations with the number of times the organizations used the sub status “Wait – User” versus the number of times the organizations were involved in any incident. This analysis was performed to identify the “organization – wait user ratio”, i.e. what is the usage percentage of status “Wait – User” by an organization in comparison with the total number of times the organization was involved in an incident. Both databases were linked based on the common data field “organization”.

We noted 10 organizations that used the sub status “Wait – User” in 5% or more of all their actions, whereby “Org line I” used it the most (in %): 20%. Note however that Org line I only used the Wait – User sub status for 2 times (i.e. 20% out of a total of 10 actions). If we compare Org line A2 with Org line C, we note that the “organization – wait user ratio” is similar (i.e. 6,77% vs. 6,60%). (See appendix Wait user behavior – 4)

5.3.5 What is the behavior per organization?

Table 24. General metrics

General metrics	*No of Records*
Number of records with sub status “Wait – User”	4.217
Number of locations	23
Number of locations who made use of the sub status “Wait – User”	20

We listed all locations with the number of times the location used the sub status “Wait – User” versus the number of times the location was involved in any incident. This analysis was performed to identify the “location – wait user ratio”, i.e. what is the usage percentage of status “Wait – User” by a location in comparison with the total number of times the location was involved in an incident. Both databases were linked based on the common data field “location”.

We noted 9 locations who used the sub status “Wait – User” in 5% or more of all their actions, whereby Germany used it the most (in %): 14,55% (i.e. 8 times Wait – User sub status used with a total of 55 actions).

We noted 9 locations who used the sub status “Wait – User” in 5% or more of all their actions, whereby Germany used it the most (in %): 14,55% (i.e. 8 times Wait – User sub status used with a total of 55 actions). (See appendix Wait user behavior – 5)

5.3.6 What is the average duration an incident is in status ‘Wait-User’?

In order to calculate how long an incident remains in the Wait-User substatus we eliminated all service requests that were still stuck in this status. When looking at the top ten of incidents that are have the longest Wait User time, we noticed that these have been in this status for more than 100 days. There even is an incident with more than 1 year of Wait-User time.

The average duration is 7,74 days. Furthermore, in 29,72% of the cases an incident is more than 1 week in this status.

Table 25. Wait user duration

SR NUMBER	*'Total Wait-User duration per SR'*
1-512795200	386,31

1-523391859	210,82
1-580987781	205,63
1-565045794	199,66
1-605313141	167,97
1-559795575	164,26
1-613379581	163,00
1-606902814	155,23
1-626766981	128,14
1-603556351	113,24

Moreover, we calculated the correlation between the the wait-user time and the throughput time of incidents: 0,54. So the waiting time has a strong impact on the total duration of the incident.

When we considered the total user waiting time in function of the total duration of all incidents together, we saw that the waiting user time contributes to the total duration time for 38,42%.

We also took a look at the portion of the waiting time in the throughput time, for Organizations A2 and C separately. For organization A2, 28,35% of their throughput time is user waiting time. For organization C it is 38,45%. (See appendix Wait user behavior – 5)

5.4 Process Conformity per Organization

First we entered the complete incident log in ProM, with our own activities. This resulted in the process flow as shown in Fig. 11. The process map shows us that the standard process flow is well followed, however a lot of transfers between support teams take place. We also see that the in a high number of cases input is requested, which points at the “Wait User” usage. It is remarkable that Register Incident is not the main starting point.

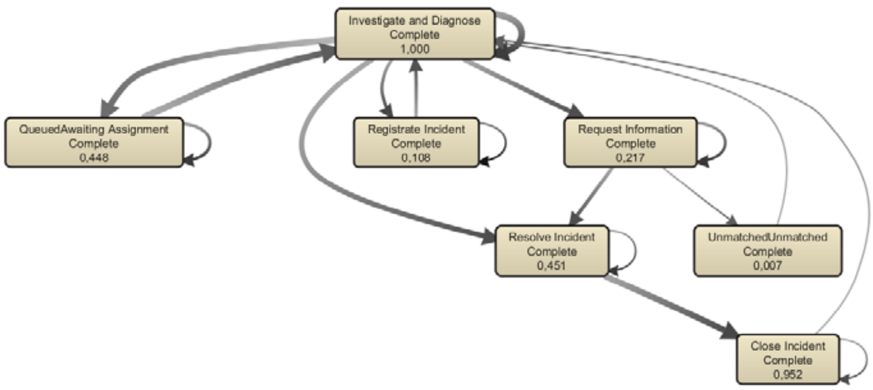

Fig. 11 Process map of the incident management process

If we compare the process flows followed in Org line A2 and Org line C, as shown in Fig. respectively Fig. ,we see a similar sequence. This shows that the process is consistently followed across the different organisations. However, from our analysis of the Push to Front process we can derive that in Org line A2 the status Queued/Awaiting Assignment is significantly more linked to a transfer to Second or Third Line.

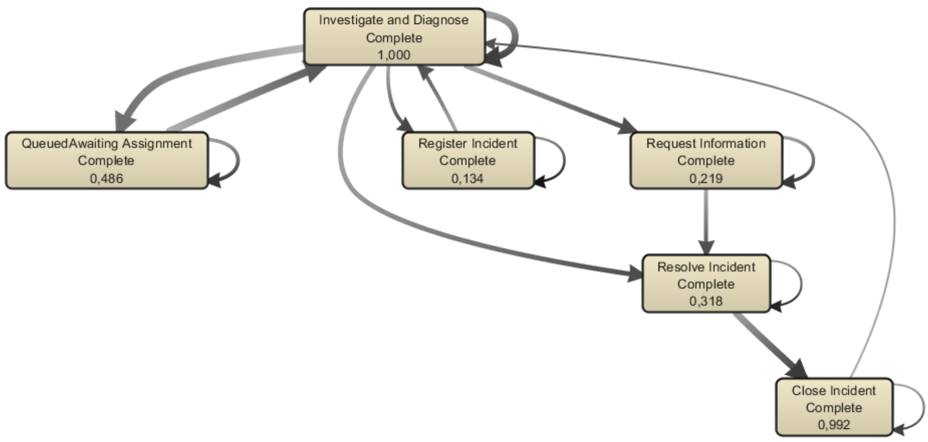

Fig. 12 Process map of the incident management process followed in Org line A2

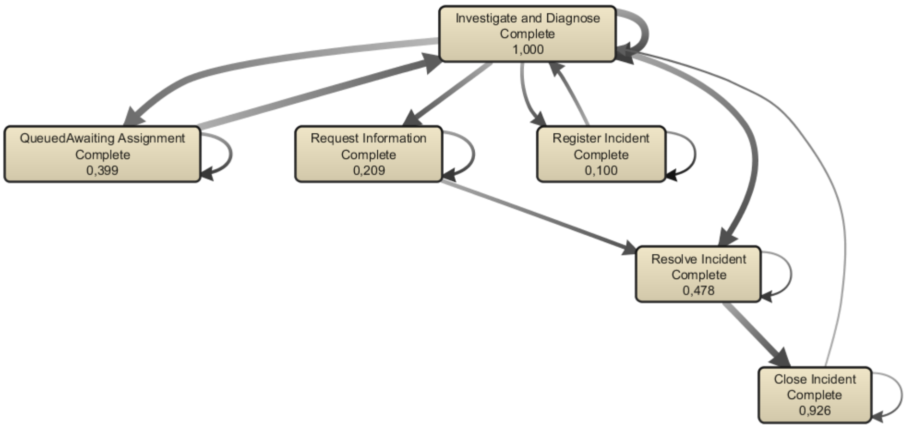

Fig. 13 Process map of the incident management process followed in Org line C

As part of the verification on whether the organisation A2 and C are operating in a similar way, we performed an comparison between both organisation on the through-put time of incidents. As can be noted, 35% of all incidents for organisation C are closed within a day. For A2 this is only 8%.

Table 26. Frequency table of through-put time (all tickets)

From	To	# Total Closed	% of Total closed	# for A2	% for A2	Cumul % for A2	# for C	% for C	Cumul % for C
0	1	2155	29%	57	**8%**	8%	1988	**35%**	35%
1	2	287	4%	48	6%	14%	176	3%	38%
2	5	286	4%	36	5%	19%	186	3%	41%
5	10	2480	33%	299	40%	59%	1776	31%	72%
10	20	1426	19%	155	21%	80%	1013	18%	89%
20	50	664	9%	87	12%	92%	471	8%	98%
50	100	150	2%	24	3%	95%	91	2%	99%
100	200	70	1%	24	3%	98%	33	1%	100%
200	500	34	0%	14	2%	100%	17	0%	100%
500	1000	1	0%	1	0%	100%	0	0%	100%
		7553	100%	745	100%		5751	100%	

Table 27. Descriptive Statistics overview

Distribution Data	*Overall*	*A2*	*C*
Mean	12,1	24,0	10,1
Standard Error	0,3	2,1	0,3
Median	7,6	8,3	7,5
Standard Deviation	28,6	57,5	22,5

Using Minitab, we confirmed that the distribution of through-put time is not normally distributed and applying the test of equal variances (i.e. Levene test) we rejected the null hypothesis of equal variances ($P < 0,05$) and concluded that there is a difference between the variances in the population.

6 Other interesting findings

6.1 When is the helpdesk called the most?

When we consider all the data over the 3 years, we can see that in general the most incidents are reported on Wednesday, Thursday and Friday.

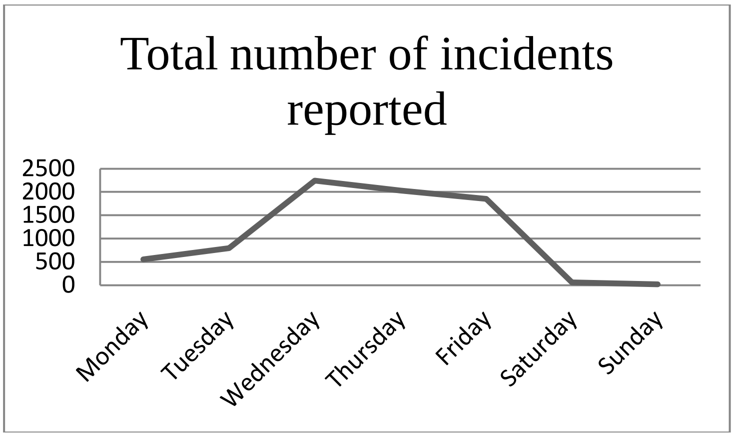

Fig.15. Reported incidents.

6.2 How long does it take to solve an incident?

We have calculated the time spans of the incidents.

Table 28. Incidents throughput time

Turnaround time of an incident	*Number of incidents*
< 1 minute	64
< 1 hour	1480
< 1 day	611
< 8 days	2229
< 15 days	1750
< 22 days	629
< 30 days	310
< 60 days	274
< 90 days	88
< 120 days	47
< 150 days	16
< 180 days	11
< 1 year	36
< 2 years	7
> 2 years	1

The average duration is 12 days, and falls thus in the category of < 15 days. Most of the incidents are solved within the week. 64 incidents are remarkably solved within one minute.

6.3 What is the most common impact of the incidents?

The majority of the incidents have a low or medium impact.

Table 29. Impact to customer

Impact	*Number of incidents*
Low	3245
Medium	4045
High	260
Major	3

The incidents with a major impact have only lasted 1 or 2 weeks. The low and high impact incidents are mostly solved within two weeks, most the medium and high impact incidents within one week. There is however one incident that has taken longer than 2 years.

Table 30. Impact versus incident throughput time

Duration	*Low*	*Medium*	*High*	*Major*	***Grand Total***
< 1 minute	34	30			**64**
< 1 hour	840	636	4		**1480**
< 1 day	299	296	16		**611**
< 8 days	638	1490	100	1	**2229**
< 15 days	801	868	79	2	**1750**
< 22 days	313	293	23		**629**
< 30 days	158	140	12		**310**
< 60 days	114	148	12		**274**
< 90 days	29	51	8		**88**
< 120 days	11	32	4		**47**
< 150 days	2	14			**16**
< 180 days	2	9			**11**
< 1 year	3	31	2		**36**
< 2 years	1	6			**7**
> 2 years		1			**1**
Grand total	**3245**	**4045**	**260**	**3**	**7553**

7 References

1. IT Governance institute: Cobit 4.1 framework, control objectives, management guidelines, maturity models (2007)
2. BSI: BS ISO/IEC 20000-1:2005 Information technology – service management
3. BSI: BS ISO/IEC 20000-2:2005 Information technology – service management.
4. Oklahoma government, OSF service support – incident management process, http://www.ok.gov/cio/documents/ServiceRequestProcessOverview.doc.

Appendix: Wait user behavior

1 Action owner – wait user ratio

OWNER_FIRST_ NAME	*NO_WAIT_USER*	*NO_INCIDENTS*	*%_WAIT_USER_VS_INCIDENTS*
Amer	25	74	33,78
MUDIT	21	66	31,82
Mohammad	27	86	31,40
Anjali	31	122	25,41
Sharath	14	58	24,14
Patryk	26	108	24,07
Muthu	80	355	22,54
Aneesh V	21	97	21,65
Prashant	27	126	21,43
Meishan	12	57	21,05

2 Support team – wait user ratio

INVOLVED_ST	*NO_WAIT_USER*	*NO_INCIDENTS*	*%_WAIT_USER_VS_INCIDENTS*
N45	19	88	21,59
L50 3rd	24	113	21,24
G73	11	55	20,00
G356 2nd	19	105	18,10
W4	11	62	17,74
V17 3rd	42	247	17,00
S24	20	118	16,95
G92	191	1154	16,55
N20	10	62	16,13
G297	15	94	15,96

3 Function – wait user ratio

INVOLVED_ST FUNCTION_DIV	*NO_WAIT_ USER*	*NO_RECORDS*	*%_WAIT_USER VS_RECORDS*
D_1	161	1488	10,82
V3_3	3	34	8,82
E_5	248	2907	8,53
E_10	383	4527	8,46
A2_1	803	9977	8,05
D_2	2	28	7,14
A2_2	180	2618	6,88
V3_2	1930	30950	6,24
E_4	31	510	6,08
E_8	15	261	5,75

4 Organization – wait user ratio

NO_ RECORDS	*INVOLVED_ ORG_LINE_3*	*%_ RECORDS*	*NO_WAIT _USER*	*%_WAIT _USER_ VS_RECORDS*	*%_WAIT _USER*
12508	Org line A2	19,09	847	6,77	20,09
4623	Org line B	7,05	355	7,68	8,42
42189	Org line C	64,38	2783	6,60	65,99
28	Org line D	0,04	2	7,14	0,05
186	Org line G2	0,28	10	5,38	0,24
861	Org line G4	1,31	12	1,39	0,28
10	Org line I	0,02	2	20,00	0,05
22	Org line V1	0,03	4	18,18	0,09
480	Org line V11	0,73	3	0,63	0,07
605	Org line V2	0,92	50	8,26	1,19

5 Location – wait user ratio

COUNTRY	*NO_WAIT _USER*	*%_WAIT _USER*	*NO_ RECORDS*	*%_ RECORDS*	*%_WAIT_USER _VS_RECORDS*
0	3	0,07	245	0,37	1,22
SE	10	0,24	544	0,83	1,84
au	7	0,17	188	0,29	3,72
be	381	9,03	5944	9,07	6,41
br	121	2,87	2660	4,06	4,55
ca	15	0,36	403	0,61	3,72
cn	83	1,97	1186	1,81	7,00
de	8	0,19	55	0,08	14,55
fr	158	3,75	3158	4,82	5,00
gb	10	0,24	267	0,41	3,75