BPIC'13: Mining an Incident Management Process

Emmy Dudok, Peter van den Brand

Perceptive Software, 22701 West 68th Terrace, Shawnee, KS 66226, United States
{Emmy.Dudok, Peter.Brand}@perceptivesoftware.com

Abstract. This paper describes the results of the exploratory process mining efforts on the incident management process event log provided by Volvo IT Belgium. Specific areas of interest provided by the process owner are analyzed as well as some additional areas of interest that qualified for further investigation based on the information provided. Interesting results include uncovering specific support teams and products for which specified unwanted behavior such as lack of push to front, ping pong, and wait user abuse was prominent. Also some interesting relations were found, e.g. between the wait user abuse and incident impact category, and the hypothesis that a correlation exists between the number of handovers and total resolution time was proven.

1 Introduction

An event log of the incident management process was provided by Volvo IT Belgium[1] for thorough analysis using business process intelligence tools. The process owner highlighted four areas of interest and posed several concrete questions regarding the processes' performance in those areas. This report attempts to provide answers to those questions.

This study qualifies as an exploratory investigation. Little information was available on strategy or goals set by the organization related to the process execution. That means that the analysis can quantify certain behaviors in the process, but it will not be possible to make any qualitative judgments about performance. Nevertheless, the findings in this report highlight patterns that are likely to be important insights for the process owner.

First, the areas of interest are discussed. Then, in the following sections, the exploratory analysis of each of those areas of the incident management process is described. Finally, the report is concluded with a summary of the findings and some recommendations.

1.1 Areas of interest

The process owner provided a total of four areas of interest: Push to Front, Ping Pong Behavior, Wait User Abuse, and Process Conformity per Organization. Additionally,

[1] doi:10.4121/500573e6-accc-4b0c-9576-aa5468b10cee

another area of interest has been added to the analysis: Resolution Verification. All are briefly discussed below.

Push To Front. Most service requests such as incidents need to be resolved by the first line support teams. If those teams cannot handle a request, they will send it to second and/or third line support teams. This should happen as little as possible though, because resolving service requests is usually not the core business of these support teams. However, as these teams have in-depth knowledge of specific products, their assistance might be required in resolving some service requests. Insight is desired into the frequency with which this behavior of "pushing back" requests to the second and third line support teams occurs. Additionally, further investigation into possible causes for the behavior is requested with a focus on differences regarding products, organizations, and functions.

Ping Pong Behavior. Ideally, service requests are resolved with as few handovers to other support teams as possible. However, there are cases in which support teams send service requests back and forth to one another. This is referred to as ping pong behavior and is a source of inefficiency in the process. Insight is requested in the frequency of occurrence of the behavior for functions, organizations, specific support teams and/or products.

Wait User Abuse. A key performance indicator for the process and measurement of support teams' and individual performance is the total resolution time of an incident. This time does not include the time spent on waiting for a response from the service request owner (the user who reported the incident). In the resolution of a service request it is possible to set the status to "Wait – User", indicating the request is put on hold until response by the service request owner is received. However, support teams can also use this status to artificially lower the total resolution time in order to make the key performance indicator look better. Insight is requested into the frequency of abuse of the "Wait – User" status for support teams and employees. Additionally, it can be expected that incidents with a higher impact will have a shorter resolution time, as they should be resolved more quickly to restore the process performance. Therefore, it might also be expected that there is a higher frequency of "Wait – User" abuse when the incident has a higher impact.

Process Conformance. Volvo IT is divided into two organizations: Org line A2 and Org line C. The process owner would like to see to what extent these two organizations conform to their standard incident management process. However, a process model of the standard process is not available, so a conformance check cannot be performed. Instead, the differences in process between the two IT organizations Org line A2 and Org line C will be investigated.

Resolution verification. When a service request is resolved it will be marked as such and sent back to the service request owner for verification. The moment the resolution is verified, the service request will be closed automatically within 24 hours. If the service request is not verified, it will be closed automatically after 7 calendar days after the status has been changed to "Resolved". As action owner, it is possible to

verify the solution and immediately close the service request. An analysis can be executed on the frequency with which resolutions to service requests are or are not being verified and how often the service request owner has verified the resolution.

2 Incident Management

2.1 Context

The incident management process can be characterized as reactive, and a formal working method is put into place in order to respond as efficiently and effectively as possible in resolving reported incidents. The incident management process starts when an incident is created. Incidents can be reported by people or are automatically generated by a system. Initially, the Help Desk or Expert Helpdesk will try to resolve the incident. If they cannot resolve the incident, it is escalated to a second or even third line support team. Incidents are resolved as soon as possible in order to restore service levels and/or resume day-to-day work. When an incident is resolved, the resolution is verified with the incident owner (reporter) and is automatically closed.

The basis for the analysis is the event log obtained from Volvo IT. It is important to know which query was used to obtain the event log in order to be able to interpret the results of the analyses. The exact details of this query are not known however. The event log seems to be based on incidents for which a status has been changed during a three-week period from the 1st of May 2012 up to and including the 23rd of May 2012. In total, a number of 7,554 incidents are present in the event log. 7,546 of these incidents are certain to have been closed as they have reached a closed status. It is important to keep in mind that the results provided in this study are only based on this specific event log. Since the event log only contains a small subset of data, results may not hold when more incidents and a larger time frame are considered.

2.2 Data Preprocessing

Some preprocessing of the data was necessary before conducting the analysis. The steps are discussed below.

Status changes can have multiple sub statuses and a sub status might be the same for different general statuses. Therefore, the status and sub status attributes were combined into the activity element.

In order to answer questions regarding push-to-front behavior, it was required to know which support tiers (first, second, and third line) an incident had passed through. In the provided event log, the support tier was part of the support team name. For example, “S2” indicated support tier 1, and “S2 2nd" indicated support tier 2. Therefore, a new attribute was created to solely indicate the support tier.

2.3 Analysis Tools

The analysis is carried out using Perceptive Process Mining, a commercial process mining tool[2]. All process models, social networks and most of the charts have been created with Perceptive Process Mining and finished in Microsoft Excel.

3 Push To Front

3.1 General

"Push to front" refers to the desire to handle as many incidents in the first line (front) support teams as possible. When incidents are being pushed back to the 2nd and 3rd tiers too often, this puts extra load on those teams. Resolving incidents is not the primary task of these teams, and the performance of the incident process suffers. The question is whether incidents are pushed back unnecessarily.

Figure 1 below shows the first part of the flow of incidents between support tiers. In total 2,156 incidents (32% of 6,702 incidents) are pushed from the 1st to 2nd and/or 3rd tiers. 4,546 incidents (68% of 6,702) started by the 1st support tier are also resolved by the 1st support tier.

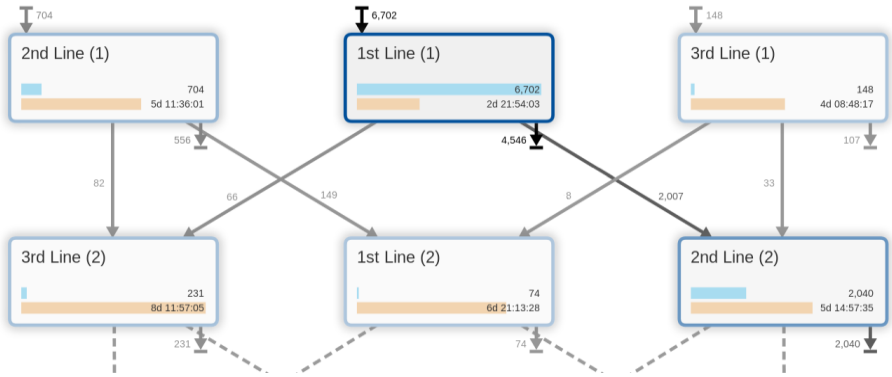

Figure 1 Routing of incidents between support tiers

Note that there are a total of 704 incidents that are not pushed to, but started in the 2nd support tier. 556 of these are handled solely by the 2nd support tier, 149 incidents are pushed to the 1st support tier and 82 incidents are pushed to the 3rd support tier. For incidents started at the 3rd support tier, 107 incidents are handled solely by the 3rd support tier, while 8 incidents are pushed to the 1st support tier and 33 incidents are

[2] http://www.perceptivesoftware.com/processmining

pushed to the 2nd support tier. Whether or not this specific behavior is desired should be discussed with the process owner.

3.2 Products

Incidents for some products might be pushed back more than for other products. In total there are 704 different products present in the event log. The analysis in this section focuses on the top 10 products with the highest number of incidents in the event log. These 10 products account for 39% of all incidents (2,594 of 6,702) started at the 1st support tier. For every product, the percentage of incidents initiated by the 1st support tier and pushed to 2nd and/or 3rd support tier is analyzed and shown in **Error! Reference source not found.**.

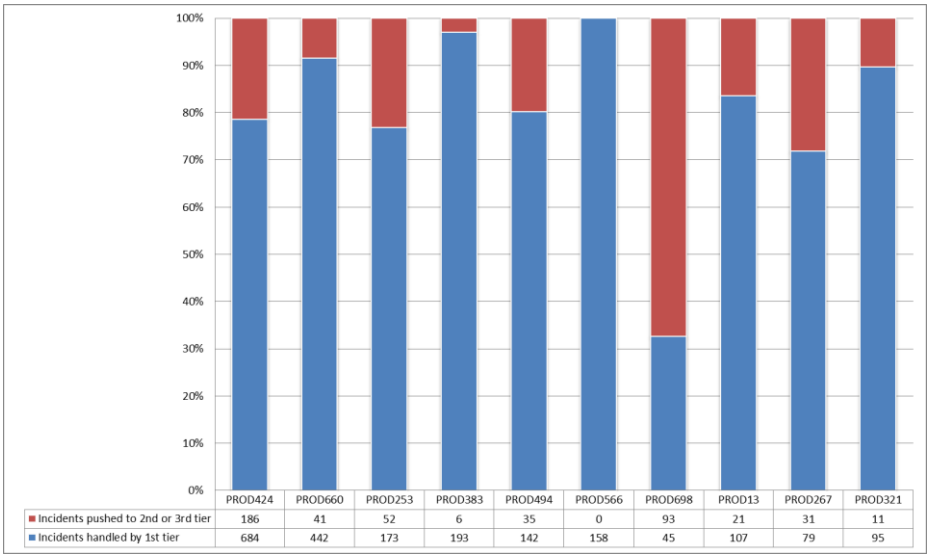

Figure 2 Incidents started at 1st tier and pushed to 2nd and/or 3rd tier, for top 10 products

Product PROD698 stands out, as 67% of its incidents (93 out of 138) are being pushed back to 2nd and/or 3rd support tiers. In contrast, PROD566 stands out because the 1st support tier handles all incidents.

Why does the 1st support line not resolve incidents for PROD698 as much as incidents for PROD566? Both products have a similar amount of incidents present in the event log. After further investigation, it is clear that incidents for PROD566 are always handled within the USA, while incidents for PROD698 are handled in 11 different countries. Furthermore, the number of people involved in the 1st support tier for PROD566 is 13, while the number of people involved in 1st support tier for PROD698 is 116. This suggests that there might be less experience with PROD698 in the first line support teams, because the incident management is not centralized for this product.

3.3 Support Teams

Some support teams might push back incidents more often than others. The percentages of incidents initiated in the 1st support tier and pushed to 2nd and/or 3rd support tier were analyzed by support team. We assume that the first support team assigned to the incident was the team that pushed the incident to the 2nd or 3rd support tiers (if it was pushed back at all). The top 10 first support teams for which most incidents are present in the event log are analyzed. They account for 77% of all incidents (5,138 of 6,702) that are started at the 1st support tier. See **Error! Reference source not found.**.

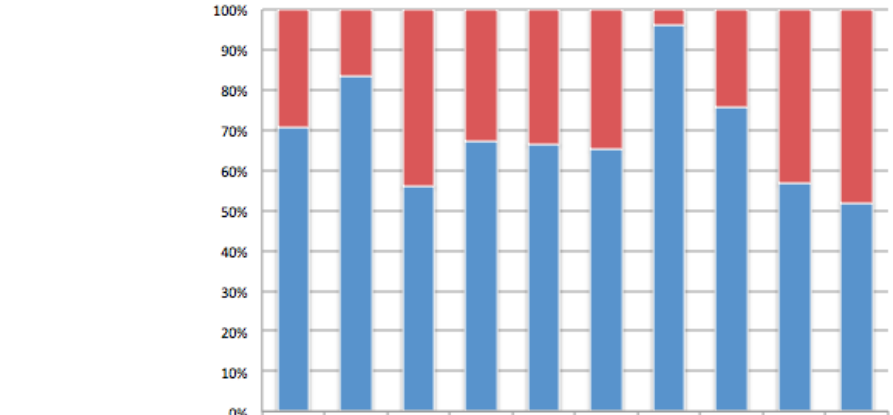

	G96	S42	G97	G76	N38	S56	N36	S37	D8	D5
Incidents pushed to 2nd and 3rd tiers	544	187	160	111	96	99	10	54	87	90
Incidents handled by 1st tier	1319	934	204	229	191	185	260	167	114	97

Figure 3 Incidents started at 1st tier and pushed to 2nd and/or 3rd tier, for top 10 first support team

Support team N36 stands out, as only 4% of all initiated incidents (10 of 270) are being pushed to 2nd and/or 3rd support tiers. On the other hand, support team G97 has a much higher percentage of incidents being pushed to 2nd and/or 3rd support tiers, namely 44% of all incidents (160 out of 364).

Why are cases initiated by support team G97 not resolved by the 1st support line as often as incidents initiated by support team N36? Both support teams have a similar amount of incidents present in the event log. Incidents initiated by support team N36 are almost exclusively handled within the USA, while incidents initiated by support team G97 are handled in 15 different countries. Furthermore, the number of people involved in support team N36 is 16, while the number of people in support team G97 is 287. Again, this might be due to specialization. Incidents might be pushed back more easily when there is less specialization.

This raises the question whether cases are being pushed back more often by support team G97 than by support team N36 because of the number of unique products for which support is provided by these teams. The hypothesis is that a team that needs to support large number of products compared to the total number of

incidents they handle would have to push back a lot of incidents because it is difficult to build up the knowledge they need to handle the incidents themselves. Incidents were initiated by support team G97 for 114 different products, while for support team N36 this is only 25. For these two specific support teams there seems to be a strong positive correlation between the number of products serviced relative to the total number of cases and the number of incidents pushed back. However, after further investigation on all support teams, this correlation does not seem to hold in general.

3.4 Impact

Impact, urgency, and priority are all indicators of the severity of a service request. When the total resolution time exceeds a specific preset period of time, the priority is increased to a higher level. It might be expected that incidents with a higher priority are pushed back more easily, because these might be more difficult to resolve and the impact on process performance is higher. Only the latest impact is available in the event log instead of the priority. Therefore, it is possible to investigate the relation between the latest impact and push to front behavior.

Figure 4 below shows the percentage of incidents pushed to 2nd or 3rd tier support teams per impact. Clearly, the higher the impact, the more often incidents are being pushed back.

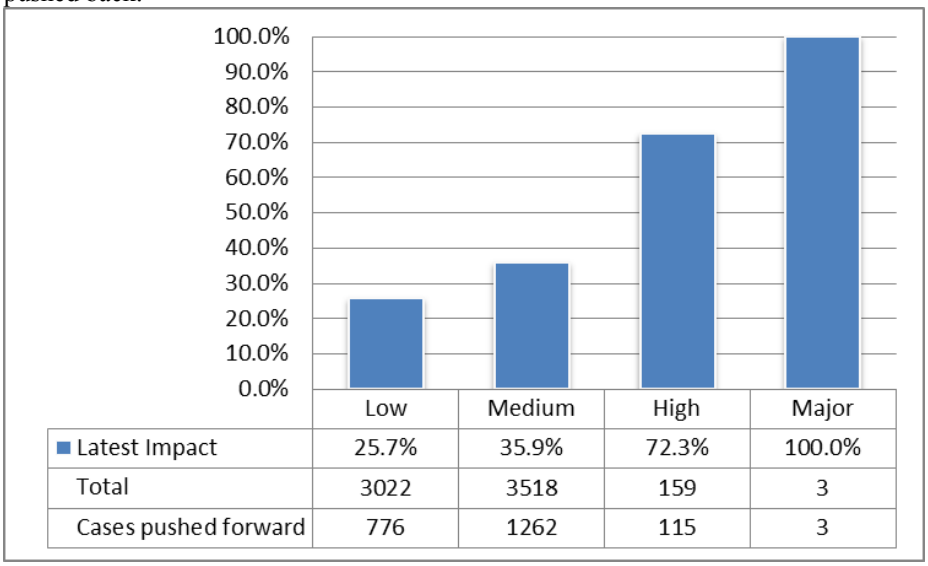

	Low	Medium	High	Major
Latest Impact	25.7%	35.9%	72.3%	100.0%
Total	3022	3518	159	3
Cases pushed forward	776	1262	115	3

Figure 4 Percentage of incidents pushed back per impact

Additionally, a negative correlation between the priority or impact and the total resolution time would be expected. The higher the impact, the lower the total resolution time should be. Figure 5 shows the average time to resolution per impact.

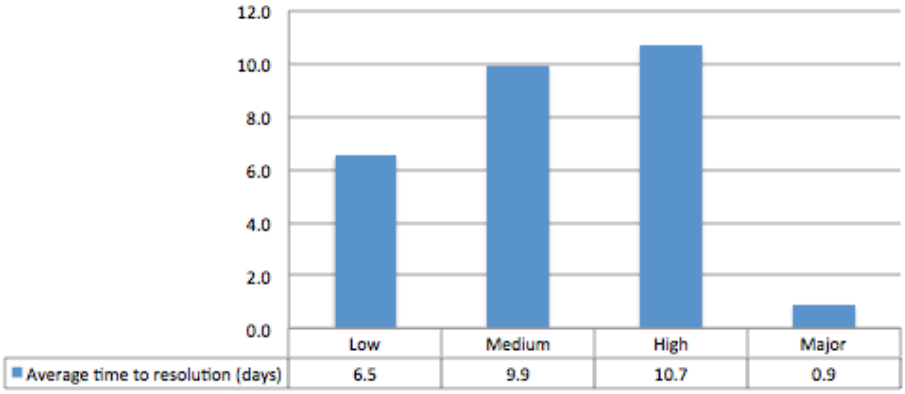

Figure 5 Average time to resolution per impact

Intrestingly, the high impact incidents have the highest average time to resolution. Major impact incidents have on average the lowest resolution time, but there are only 3 major impact incidents.

Conclusion. Although it is possible to indicate whether the behavior of cases being pushed back occurs, it is difficult to indicate whether the observed occurrences are being pushed back legitimately or not without any inside knowledge on the process. Based on previous experience with incident management processes, pushing 30% of all incidents to higher support tiers is considered to be normal. Some products were pushed back more easily than others. To a large extent, the variations in push-to-back behavior can be explained by a lack of specialization in some of the 1st tier support teams. More specialization would likely reduce unwanted push-to-back behavior.

Another valuable insight would be to analyze whether incidents are pushed forward to 2nd and 3rd support tiers more often in busy periods. This would require insight in all open and closed incidents over a longer period of time. Unfortunately, the data was insufficient to analyze this.

4 Ping Pong Behavior

4.1 General

Ping pong behavior refers to sending incidents back and forth between support teams before they are resolved. Unnecessary ping pong behavior is a source of inefficiency in the process, and is a possible cause for not satisfying the service level agreements that are in place. In this section it is analyzed where this behavior is most prevalent.

There are two ways to indicate the handover (ping pong) behavior. The first is to look at the frequency with which a particular status change has occurred. In this case, the status “Queued – Awaiting Assignment”, which is used when a service request is

assigned to someone else, can be regarded as the initiation of the handover of work. However, it might be that this assignment is not always executed and work is handed over more often to other support teams or less often when it is handed over within a particular support team than would appear when only considering this status change. Therefore, another, more accurate option is to actually consider the number of times the incident has been handed over to another support team, i.e. when another support team has changed a status for the incident.

We distinguish two types of ping pong behavior: "linear" and "circular". Both are characterized by having a high number of handovers of work between support teams for a specific incident. In linear ping pong behavior the case is going back and forth between just a very small number of teams. Circular ping pong on the other hand is characterized by incidents continuously being pushed to other support teams.

The distinction is made because the hypothesis for the root cause is different for each. Based on experience with other incident management processes, it is expected that the most important cause for circular ping pong is the fact that incidents get routed to the wrong support team because of incorrect or incomplete information about the incident. The support team receiving such an incident concludes they are not the right team to solve it, and then makes a best guess based on that information when routing it to another support team. Of course, this pattern repeats itself until the incident finally gets routed to the correct team, or when more information becomes available about the incident, which allows correct routing. In contrast, a common cause for linear ping pong cases is disagreement over who should handle a particular incident.

Circular ping pong is distinguished from linear ping pong by looking at the number of different support teams that are involved. When 4 or more support teams are involved, an incident is considered to exhibit the circular ping pong behavior. Conversely, when only 2 or 3 support teams are involved then an incident is considered to have linear ping pong behavior.

To qualify as either type of ping pong behavior, an incident requires a high number of handovers of work. When there are at least 4 handovers from one support team to another, unwanted ping pong behavior is assumed to have taken place. The first and last handover are considered to be part of the normal process if a 2nd or 3rd tier support team is involved. That means that the criterion of 4 handovers selects incidents with at least 2 unnecessary handovers. Although both could have valid reasons for occurrence, the overall frequency of occurrence should be very low.

4.2 Lifetime of an Incident

Volvo IT Belgium suggests that there is a correlation between the lifetime (throughput time) of an incident and the amount of ping pong behavior of that incident. Figure 6 below shows the average resolution time in days of incidents with 0 to 9 handovers, which accounts for 99% of all incidents (7,465 of 7,554). A handover is defined as routing a case from one support team to another. Incidents with up to 31 handovers were observed in the event log, but the volumes of incidents with 10 or more handovers (the other 1% of incidents) are too small to discover a meaningful

pattern. Clearly, the more handovers the higher the throughput time in this process. The question remains what the possible causes might be though.

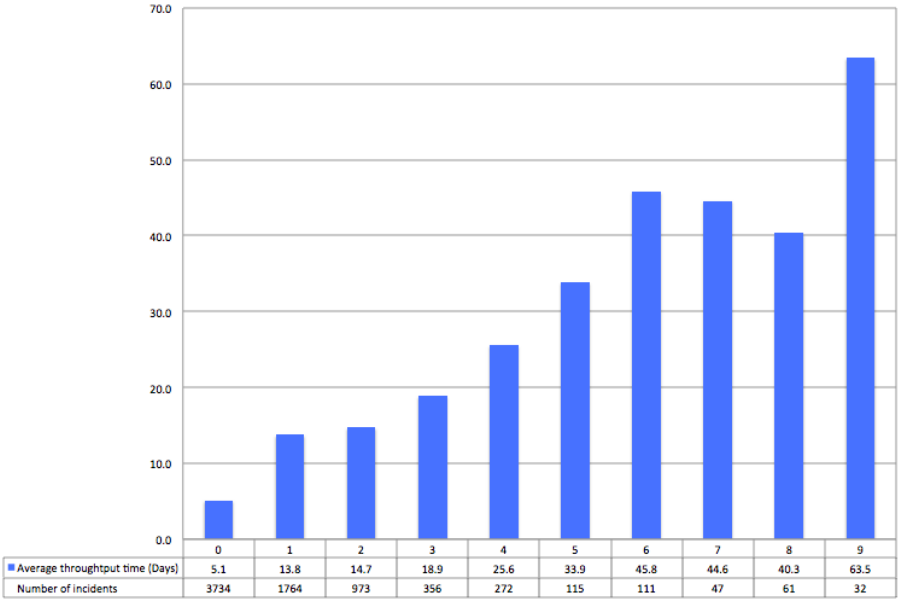

	0	1	2	3	4	5	6	7	8	9
Average throughtput time (Days)	5.1	13.8	14.7	18.9	25.6	33.9	45.8	44.6	40.3	63.5
Number of incidents	3734	1764	973	356	272	115	111	47	61	32

Figure 6 High number of handovers mean high resolution times

4.3 Circular and Linear Ping Pong

As mentioned before, we make a distinction between circular and linear ping pong behavior. In this section, only incidents are considered which were initiated in the 1st support tier (a first line support team). 9% of all incidents (670 of 7,554) exhibit ping pong behavior. 53% (354 out of 670) incidents have circular ping pong behavior, while 47% (316 out of 670) incidents are linear according to our definition above.

Figure 7 below shows an extreme case of circular ping pong, where a single incident was routed via 8 different support teams in several ‘circles’.

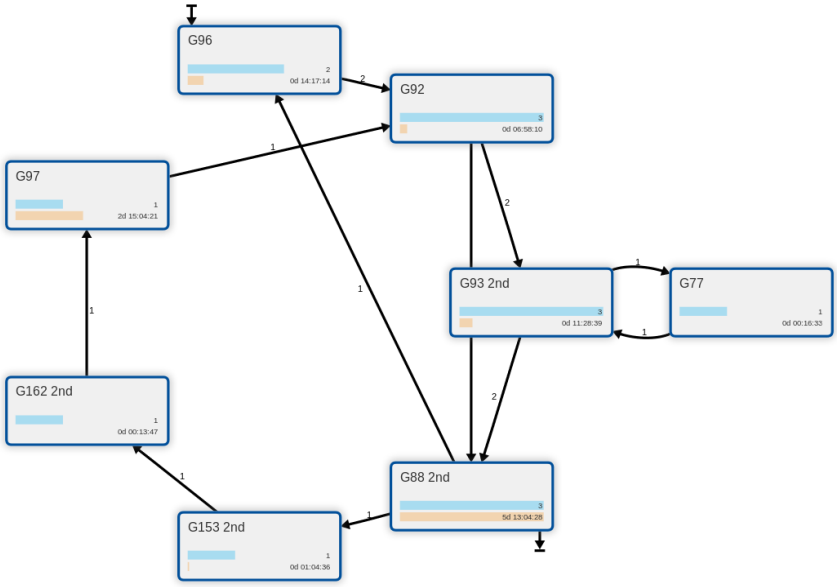

Figure 7 Single incident with extreme circular ping pong behavior

Figure 8 below shows an extreme case of linear ping pong, where a single case was handed over 20 times, of which 16 times between just two support teams. Anecdotally, in the incident shown in Figure 10 (Case ID 1-506071646), team D8 first receives the incident and reassigns it to team V37 2nd. Team V37 2nd in its turn routes it back to D8 within 2 hours. This cycle repeats 8 times. This pattern might exemplify a situation in which teams only 'communicate' by reassigning tickets to each other without any other form of communication. The result is a large delay in actually resolving the incident: the incident shown in Figure 8 took over a year to resolve.

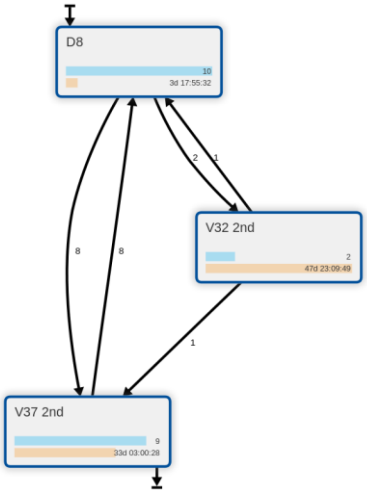

Figure 8 Single incident with extreme linear ping pong behavior

Figures 9 and 10 below show the top 10 support teams with respect to the number of ping pong incidents they were involved in. Half of the top 10 teams appear in both charts, meaning that these teams are involved in both types of ping pong. The differences among these teams are still significant though, with team D4 being involved in around 50% less circular and 50% more linear ping pong incidents than team G97[3].

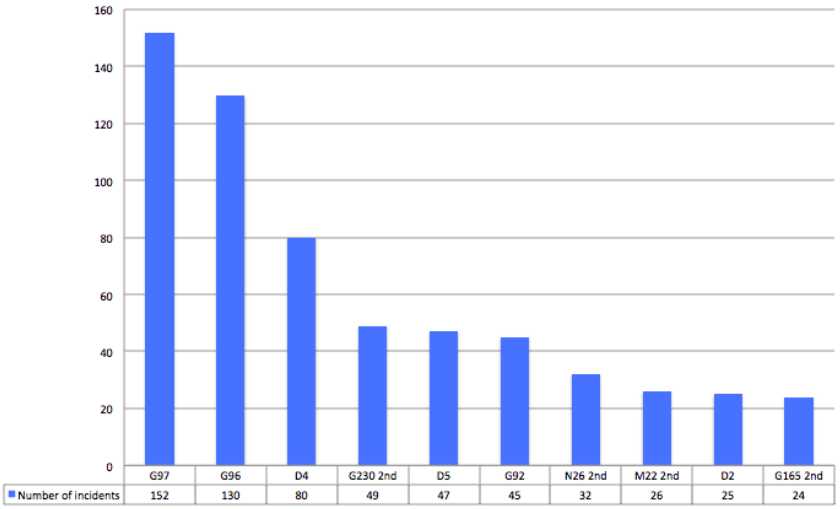

Figure 9 Top 10 support teams involved in circular ping pong incidents

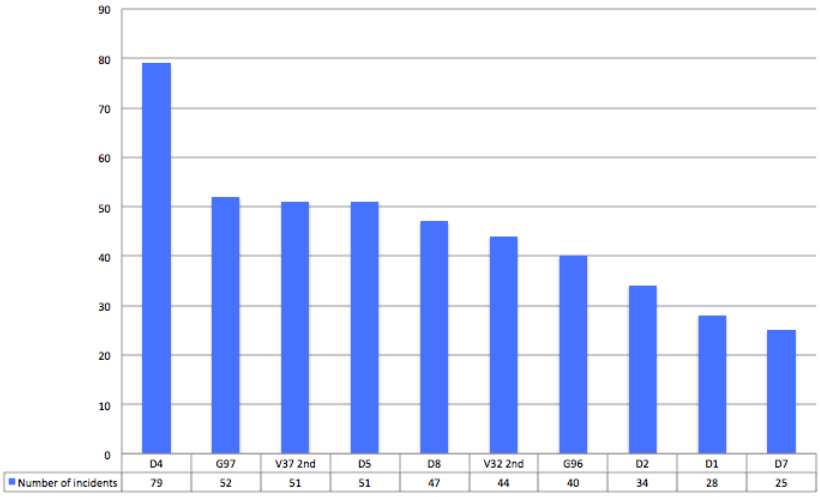

Figure 10 Top 10 support teams involved in linear ping pong behavior

3 It should be noted that D4 is involved in a total of 208 incidents, while G97 is involved in a total of 964 incidents. We chose to present the analysis based on absolute numbers, because we want to identify the major sources of handovers.

When the social networks of support teams D4 and G97 are considered, the difference becomes very clear as well. D4 has more linear behavior with incidents routed back and forth to the same other support team, whereas support team G97 is involved in more circular behavior with incidents being routed to different support teams over and over again. Moreover, the social network of D4 has fewer nodes and therefore hands over work to fewer support teams than support team G97.

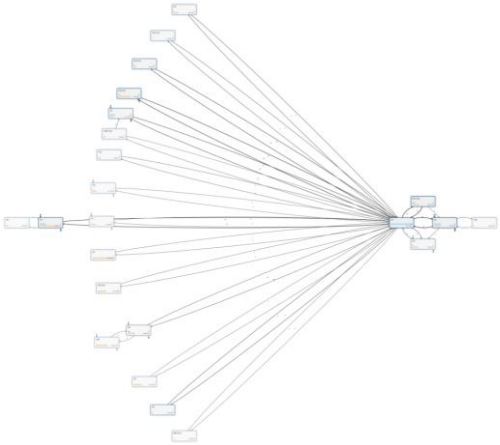

Figure 11 Social network of support teams for incidents involving D4

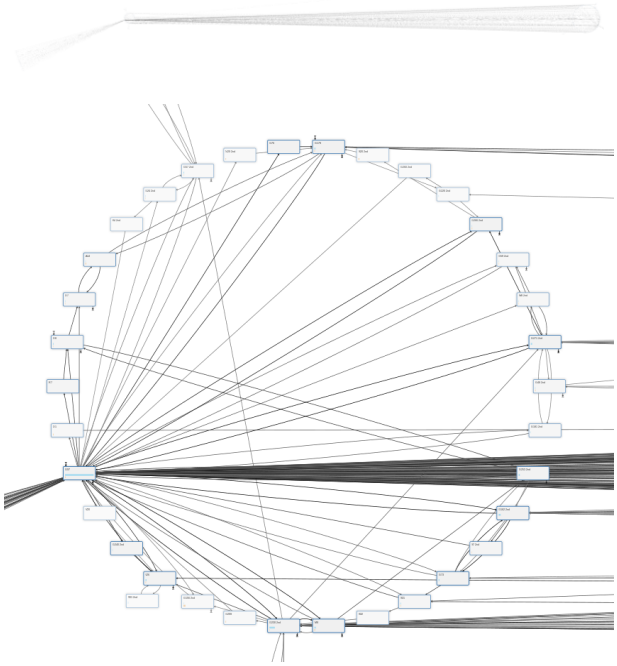

Figure 12 Social networks of support teams for incidents involving G97

4.4 Ping Pong for Products

Ping pong behavior might occur more often for some products than for others. In the Figure 13 below, the average number of handovers per product is shown, sorted by number of incidents. Product PROD542 stands out as it has by far the highest average number of handovers. Further investigation shows that all incidents for this product have been handed over at least 4 times, see Figure 14.

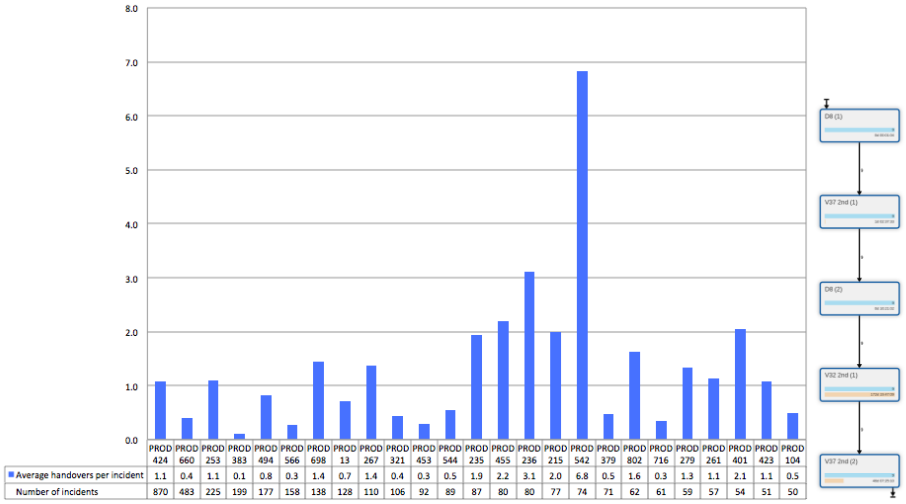

	PROD 424	PROD 660	PROD 253	PROD 383	PROD 494	PROD 566	PROD 698	PROD 13	PROD 267	PROD 321	PROD 453	PROD 544	PROD 235	PROD 455	PROD 236	PROD 215	PROD 542	PROD 379	PROD 802	PROD 716	PROD 279	PROD 261	PROD 401	PROD 423	PROD 104
Average handovers per incident	1.1	0.4	1.1	0.1	0.8	0.3	1.4	0.7	1.4	0.4	0.3	0.5	1.9	2.2	3.1	2.0	6.8	0.5	1.6	0.3	1.3	1.1	2.1	1.1	0.5
Number of incidents	870	483	225	199	177	158	138	128	110	106	92	89	87	80	80	77	74	71	62	61	59	57	54	51	50

Figure 13 Average number of handovers per product

Figure 14 Most frequent path for PROD542

Conclusion. A total of 9% of all incidents are considered to have ping pong behavior. 53% of these were considered circular ping pong, whereas 47% was considered linear ping pong. The hypothesis that a correlation exists between the number of handovers and total resolution time was also proven. A difference for the support teams was observed for the frequency with which they are involved in ping pong behavior. Moreover, a difference was found between teams and products more involved in circular or linear behavior, or even highly involved in both.

5 Wait User Abuse

5.1 General

Wait user abuse is said to occur when the status is set to “Wait – User” while no response by the owner is actually required. This distorts the actual resolution time of

the incident. In this part an analysis is performed on the wait user abuse and where the use of this status is most prevalent.

5.2 Impact on Resolution Time

Whether or not the use of this status was legitimate is quite difficult to assess, as no additional information is present in the event log nor in the additional information provided. However, as this behavior impacts the KPI of the resolution time, we are able to analyze the impact the wait user abuse has on achieving this KPI.

The only closely related information available from the VINST user manual is the time before escalation to another priority. The exact priorities and urgencies are not available in the event log, so we will approximate KPIs using the Latest Impact data field in the event log. We assume medium urgency for all incidents, and we assume that an incident with a particular impact value should be resolved within the time before increase from medium to high urgency. For example, the KPI for Low impact incidents is assumed to be 1 week, and the KPI for major impact incidents 2 hours. See Table 1 below.

Impact	Low			Medium			High			Major		
Urgency	Low	Med	High	Low	Med	High	Low	Med	High	Low	Med	High
Priority	12	11	10	9	8	7	6	5	4	3	2	1
Base	3w	1w	3d	1w	3d	1d	3d	1d	4h	4h	2h	

Table 1Escalation Matrix from VINST Manual

As is usual for incident management processes, we assume the resolution time is measured until the moment the resolution is provided, and not the moment on which the incident is closed. The charts below (see Figure 15) show histograms of the resolution times of all closed cases specified by their impact. Horizontal lines are shown for each of the three possible urgency levels.

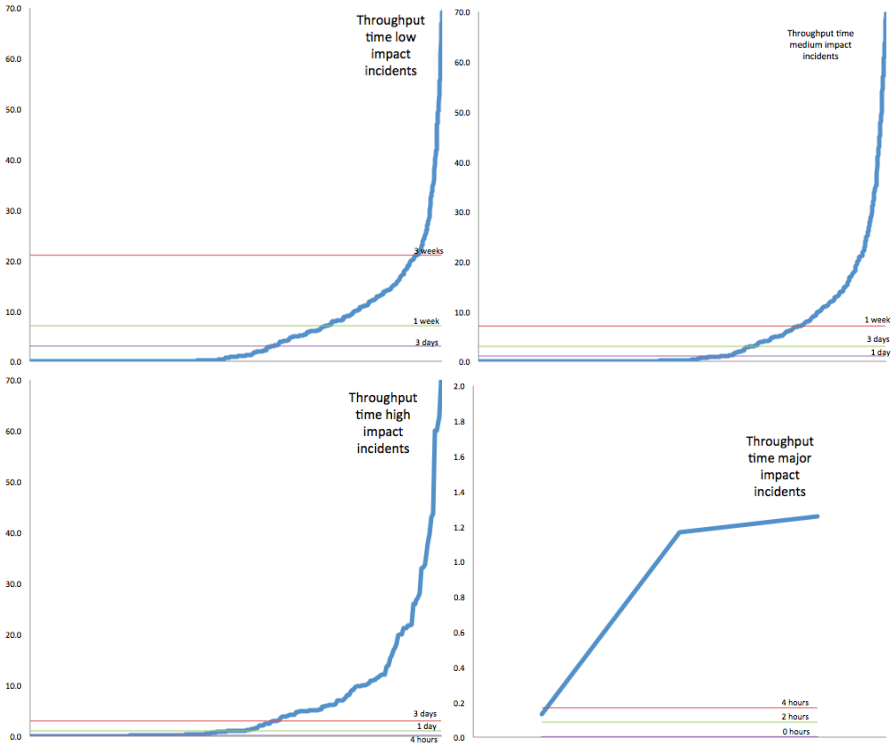

Figure 15 Histogram of resolution times per impact

Assuming medium urgency, which is the green horizontal line in the charts above, the impact of "Wait – User" abuse has on achieving the KPIs can be investigated.

Latest Impact Medium Urgency KPI	% Within KPI Total		% Within KPI Without "Wait-User"	
Low	71.3%	(2,310 of 3,242)	81.6%	(2,644 of 3,242)
Medium	65.3%	(2,639 of 4,042)	77.1%	(3,117 of 4,042)
High	47.1%	(122 of 259)	59.1%	(153 of 259)
Major	0%	(0 of 3)	0%	(0 of 3)
Total	67.2%	(5,071 of 7,546)	78.4%	(5,914 of 7,546)

Table 2 Performance on Latest Impact KPI

In total, 67.2% of all closed incidents are resolved within the KPI when the total throughput time is considered. When the time is subtracted for which incidents have resided in the status "Wait – User" a total increase of 11.2% more incidents that are resolved within the KPI are observed.

The rest of this section focuses on these 11.2% of incidents, because these are the incidents for which usage of the "Wait – User" status allowed the support team to resolve the incident within the KPI. This is exactly the set of incidents where

(successful) abuse of the “Wait – User” status would occur. We label these incidents as “possible abuse” incidents.

Figure 16 below shows the time spent in the “Wait – User” status for “possible abuse” incidents versus other incidents, split by impact. The number of incidents in the table under the figure is the number of incidents for which there was at least one occurrence of “Wait – User”.

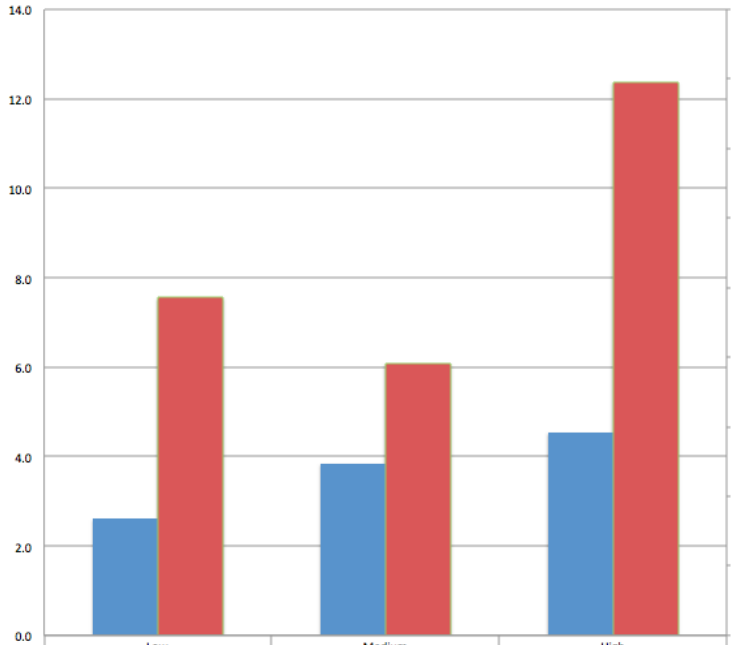

	Low	Medium	High
Average time (no abuse)	2.6	3.9	4.5
Average time (potential abuse)	7.6	6.1	12.4
Number of incidents (no abuse)	852	744	51
Number of incidents (potential abuse)	334	478	31

Figure 16 Average time spent in "Wait-User" per impact

The hypothesis is that people reporting an incident (the user that is being waited on) respond, on average, in the same amount of time regardless of whether the incident will make the KPI or not. However, the average time spent in the “Wait – User” status[4] is significantly higher for incidents that achieve the KPI after subtracting the time spend in the “Wait – User” status (the possible abuse incidents). So the hypothesis clearly does not hold. This is a strong indication that support teams are not just waiting for users to respond. Abuse of the status to achieve the KPI looks very plausible.

[4] We did not remove outliers in the averages shown in the chart. The pattern only gets more pronounced if we do remove outliers though.

5.3 Possible Causes

For the 843 incidents that are not handled within the KPI when the total resolution time is considered, but are handled when the time in the status "Wait – User" is subtracted, it can now be investigated further whether this considers specific support teams. Figure 17 below shows the percentage of "possible abuse" incidents relative to the total number of incidents that achieved the KPI (including the possible abuse incidents) for the top 10 support teams that handled most incidents.

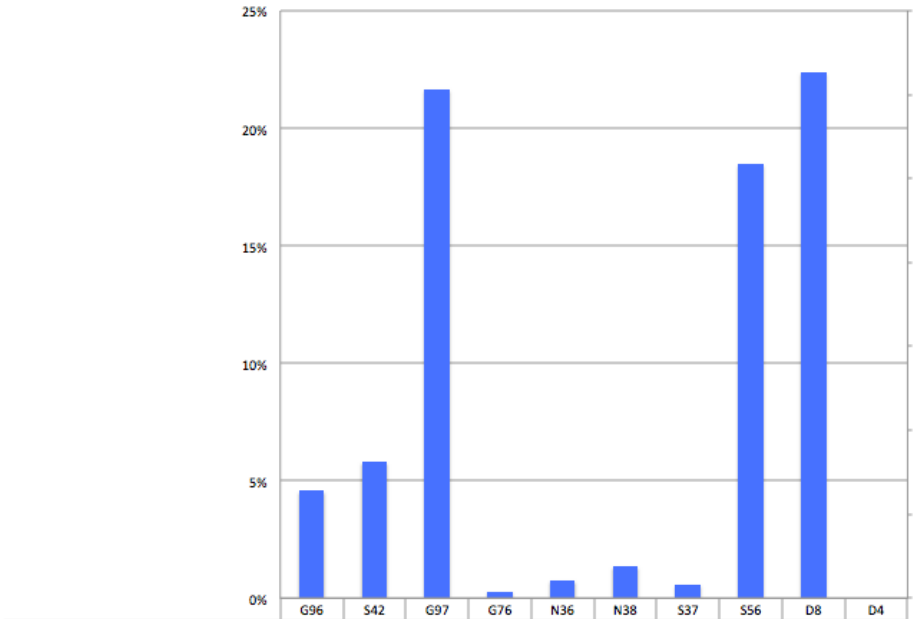

	G96	S42	G97	G76	N36	N38	S37	S56	D8	D4
Percentage of incidents with possible abuse	5%	6%	22%	0%	1%	1%	1%	18%	22%	0%
Total number of incidents within SLA	1275	958	398	350	264	219	173	128	111	56
Number of incidents with possible abuse	61	59	110	1	2	3	1	29	32	0

Figure 17 Percentage of "possible abuse" incidents relative to number of incidents that achieved KPI

Support teams G97, S56, and D8 have a high percentage of incidents for which they potentially abused the "Wait – User" status. Support teams with a high percentage could be further investigated with the process owner.

Conclusion. Clearly, the use of the status "Wait – User", legitimate or not, highly impacts the performance measurement of the total resolution time. A total of 11.2% of all closed cases achieve the Latest Impact KPI when wait user time was deducted. Therefore, these incidents possibly exhibit wait user abuse. Interestingly, the average time spent in the wait user status was significantly higher for incidents that were prone to the wait user abuse than incidents that were not. This is a strong indication that support teams are not just waiting for users to respond. Additionally, some support teams had a high percentage of incidents for which they potentially abused the "Wait – User" status.

6 Process Conformance

6.1 General

When a reference model is available for the process, it is possible to check the process conformance of the execution of the process. As the reference model of the incident VISITS Support Process is not available, it is not possible to conduct a process conformance check on the reference model. However, it is possible to compare the process execution models for the two IT organizations Org line A2 and Org line C.

6.2 Org line C versus Org line A2

In total, Org line C is involved in resolving 6,120 incidents whereas Org line A2 is involved in the resolution of only 1,795 incidents[5]. Org line C has used two more sub statuses than Org line A2: "Completed – Cancelled" and "Unmatched – Unmatched". In Figure 18 below, the most frequent paths of the organizations are compared. A blue edge surrounding an activity represents that it is performed by both organizations; orange represents that it has been performed only by Org line C, and a dotted line means it has been performed only by Org line A. Notice that Org line C resolves most incidents (27.8%, or 1,698 of 6,120) in call.

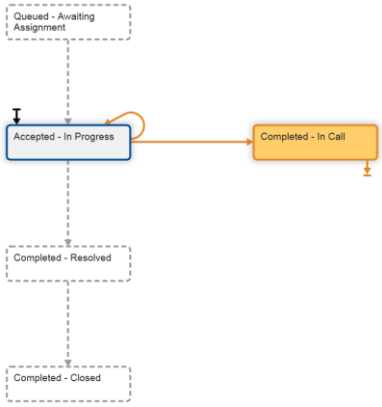

Figure 18 Comparison of most frequent path of Org line A2 and Org line C

There is a significant difference between the two organisations with respect to the frequency of pushing incidents to the second and third line. For Org line C, 29.1% of all incidents (1,625 of 5,581) that are started in the first support tier by Org line C are pushed to the second and third support tier. For Org line A2 this is figure is 54.3% (171 of 315).

[5] Only closed cases are considered in this analysis.

6.1 Conclusion

The organizations differ mostly in the number of incidents handled in this particular period. Furthermore, the most frequent paths differ in that Org line C handles most cases within first line support teams, whereas Org line A2 does not. Also, a significant difference was found between the two organisations with respect to the frequency of pushing incidents to the second and third line. When the reference model of the incident VISITS Support Process[6] would have also been available, the two process instances of the organizations of interest could have been compared to it in order to check the conformance with the reference model.

7 Resolution Verification

7.1 General

When an incident is resolved, the resolution should be verified by the service request owner. The last action owner however, can also verify the resolution without having it verified by the original owner. It might be interesting to see how often incidents are being verified, and if so how often this happens by the original owner. If incidents are not verified they are automatically closed within 7 calendar days after the status has been changed to 'resolved'. If the resolution is verified, the incident is automatically closed within 24 hours. As there is no status present in the log file to indicate the verification of the resolution, it cannot be investigated whether the resolution has been verified by the service request owner (although the service request owner can be deducted from the event log).

7.1 Results

An indication of the frequency with which incidents are being verified can be provided by having a look at the time between the status "Completed – Resolved" (first occurrence) and "Completed – Closed" (last occurrence). When a maximum of 7 calendar days has been between the two statuses, this is an indication of the incident having been verified. When more than 7 calendar days have been between the resolution and the closure of the case, this is an indication of the incident not having been verified at all. In total, 5,573 incidents are resolved and later closed. Of these incidents, 33.1% (1,846 of 5,573) is closed within 7 calendar days after the incident has been resolved. This means that it is likely that these incidents have been verified. 66.9% (3,727 of 5,573) is closed after 7 calendar days. This means that for these incidents, it is likely that no verification has taken place.

[6] VINST USER GUIDE, version 3.13 08/09/2011, P.58, http://www.win.tue.nl/bpi2013/lib/exe/fetch.php?media=vinst_manual.pdf

Conclusion. Interestingly, for 66.9% of the considered incidents it seems that no verification has taken place at all. When verification is considered part of the standard process, this percentage is rather high. Two things might explain this behavior. First, incidents can be reported automatically by the system, making verification unnecessary. However, this was the case for only 471 incidents. Reasons for this behavior should be further investigated.

8 Conclusion

8.1 Overview of results

Here, an overview of the results from the exploratory analysis is provided. No inside knowledge on the process let alone the specific products and their specifications was available to interpret the results. This makes it a challenge to indicate actual causes for the observed behavior. Some possible causes are indicated however, which could serve as a guideline for further analysis.

Results from the push-to-back behavior analysis indicate that about 30% of all incidents are pushed to the 2nd and 3rd support tier by the 1st support tier. Based on previous experience with incident management processes, pushing 30% of all incidents to higher support tiers is considered to be normal. It is highly recommended that the relation between the push-to-back behavior and busy periods is investigated. This might result in a change in resource planning to better cope with these periods. Some products were pushed back more easily than others. This might well be caused by the fact that the resolution of incidents for this product is not centralized. Further investigation is required to verify this. If there is a significant correlation, centralization of products might be worthwhile to consider. There also seemed to be a relation between the number of products a particular support team serviced and the frequency of push-to-back behavior.

Results from the ping pong analysis indicate that 9% of all incidents are considered to have ping pong behavior according to our definition. 53% of these were considered circular ping pong, whereas 47% was considered linear ping pong. The hypothesis that a correlation exists between the number of handovers and total resolution time was also proven. A difference for the support teams was observed for the frequency with which they are involved in ping pong behavior. Moreover, a difference was also found between teams more involved in circular or linear behavior, or even highly involved in both. Both teams G97 and D4 could be further investigated as they were highly involved in this behavior. In such an analysis, causes might be found that could also hold as explanation for occurrence of this behavior in other teams. Additionally, teams that are not highly involved in this behavior could be investigated to find possible reasons for their absence in this respect. Another interesting result found was that particular products, e.g. PROD542, had a high frequency of being involved in ping pong behavior in general. Further analysis should show possible causes, e.g. a lack of information on the incident, for this behavior in particular products.

Results from the wait user abuse analysis indicate that 11.2% of all closed incidents possibly exhibit this behavior according to our definition. Impact analysis on resolution times shows that 67.2% of all incidents achieve the Latest Impact KPI when wait user times are included. When the wait user time is subtracted, another 11.2% of all incidents achieve the defined KPI. Therefore, these incidents are prone to the wait user abuse behavior and were further investigated. Interestingly, the average time spent in the wait user status was significantly higher for incidents that were prone to the wait user abuse than incidents that were not. This is a strong indication that support teams are not just waiting for users to respond. Additionally, some support teams had a high percentage of incidents for which they potentially abused the "Wait – User" status. These teams might be further investigated as to whether or not they legitimately used the status.

The process conformance analysis between Org line A2 and Org line C shows that the way incidents are mainly resolved is different. Org line C has handled more incidents, and usually in first call resolution in contrast to Org line A2. Additionally, a significant difference was found between the two organisations with respect to the frequency of pushing incidents to the second and third line.

The resolution verification analysis shows that for 66.9% of all resolved and closed incidents, no resolution verification has taken place. It is recommended to reconsider the added value of this verification to the process and thereby performing an impact analysis on omitting this particular activity from the process.

In general, it is recommended to repeat the performed analysis on an event log from a longer period of time in order to verify whether the results still hold. When this is available, demand distributions can also be analyzed. Moreover, in the next paragraph some suggestions are made to add additional data to the event log for further analysis.

8.2 Additional analysis opportunities

Additional information. Inspired by the VINST Manual some additional recommendations are made to investigate some more areas of interest. In order to be able to investigate these areas, some additional information is required.

When the actual impact, urgency and priority would be available for each event, more elaborate and precise analysis could be performed regarding impact analysis.

The frequency with which resolutions are verified could be further investigated when additional information, e.g. the service request owner and resolution verifier, on the actual verification would be available.

Insight in the frequency with which actual solutions are provided following the resolution of an incident could be worthwhile to investigate. The creation of an actual (software) solution initiates another process and thus entails more work. The handover of work between the two processes might be prone to ping pong behavior for example.

The VINST Manual suggests the possibility to use templates. Templates can be used to enforce activities to resolve service requests. Because the work is standardized in a template, they could be used to reduce resolution time. When templates are used,

it could be investigated whether this actually reduces work and whether these templates are used efficiently.

Based on the service requests substatus, the system (VINST) calculates a suggested time for work effort. When the system is used to signal and support the activity owner it is recommended that the actuals for the work effort are compared to the estimated time to see whether the estimates are valid or have to be reconsidered.