

Process Mining in Information Technology Incident Management: A Case Study at Volvo Belgium

Arjel D. Bautista, Syed M. Kumail Akbar, Anthony Alvarez, Tom Metzger, Marshall Louis Reaves

CKM Advisors, 711 Third Avenue Suite 1806, New York, NY, USA
{abautista, sakbar, tmetzger, aalvarez,
mlreaves}@ckmadvisors.com

Abstract. The goal of this study is to identify opportunities that improve operational performance of information technology incident management at Volvo, Belgium. Findings are derived exclusively from computational analysis of incident and problem event logs (totaling 74,544 events) from May-June 2012, provided as part of the 2013 Business Processing Intelligence Challenge. Improvements that increase resource efficiency and reduce incident resolution times and subsequently customer impacts were identified across the following areas: service level push-to-front, ping pong between support teams, and *Wait-User* status abuse. Specific products, support teams, organizational structures, and process elements most appropriate for further study are identified and specific analyses are recommended. We conclude that operational improvement areas can be elucidated exclusively from obfuscated event logs.

1 Introduction

Incident management has attracted growing attention from process mining practitioners seeking to identify efficiency opportunities within complex business functions in recent years. Already, several studies have demonstrated the value of process mining within incident management for the purposes of improving compliance and managing risk [1,2]. Our aim is to investigate incident management with the specific objective of improving operational performance and increasing productivity. The 2013 Business Processing Intelligence Challenge (BPIC 2013) is one such opportunity to uncover sources of performance improvement in incident management by analyzing a set of real world data.

1.1 Approach and Scope

The BPIC 2013 focuses on the incident and problem management procedures of Volvo IT Belgium, from which a body of data has been collected. In our analysis of this information, we sought to understand the Volvo IT service protocols in great detail and at varying levels of granularity. In doing so, we combined the use of

process mining and computational tools with traditional spreadsheet modeling techniques to generate meaningful insights from the provided data sources.

2 Materials and Methods

2.1 Description of the Data

The event log consists of three sections obtained from Volvo IT Belgium. *VINST cases incidents* concerns the organization’s incident management segment, while *VINST cases open problems* and *VINST closed problems* contain data for the problem management system. Problems are defined as those incidents carrying a “major” impact at any point in the resolution process, or incidents that could possibly recur in the future (as judged by action owners) [3]. All three data sets contain information for cases resolved in May 2012 (with a limited number of exceptions). Each of the data sets contains analogous fields that reveal key information about steps being performed throughout the lifetime of a case [3].

Event Log	# Events	# Distinct Cases
VINST cases incidents	65,533	7,554
VINST cases open problems	2,351	819
VINST cases closed problems	6,660	1,487
Grand Total	74,544	9,860

Table 1: Quantification of events and distinct cases in each of the three data sets.

The bulk of our efforts were spent on analyzing the incident data set. We chose to prioritize our analysis and focus on incidents because they represent a majority of all cases. This enabled us to segment the data further and arrive at more pointed analysis and recommendations. Furthermore, problems typically require more tailored responses than incidents, thus inhibiting our ability to draw meaningful conclusions that can become broadly applicable business recommendations without accessing more data about the nature of the problems being resolved. Therefore, we focus our analysis on incidents.

2.2 Software Used for Analysis

We procured the version of Disco made available for the purpose of this competition (Version 1.3.6; Fluxicon, Eindhoven, The Netherlands) and loaded a project set created specifically for the BPIC 2013 original XES / MXML files [4]. We used this tool to classify cases according to path and sequence qualities difficult to represent in tabular form.

We used Microsoft Excel (Microsoft Office 2010; Microsoft Corporation, Redmond, WA, USA) in processing of the raw event logs and to explore processed data. Excel was especially helpful for performing basic and intermediate mathematical functions.

We leveraged the R software (version 3.01) with RStudio (version 0.97.449) environment for its statistical and graphical capabilities. We found both built-in and user-defined functions invaluable for preparing, analyzing, and visualizing data.

3 Data Preprocessing

3.1 Making Sense of the Raw Event Log

The BPIC 2013 data set required preprocessing prior to use in analysis and generation of meaningful business insights. This data set also posed unique problems due to the level of abstraction. Below we describe some of the cleanup and processing steps we performed and the assumptions made during our analysis of the data.

Unique Mapping of Action Owners

The only name field given, *Owner First Name* (1,440 unique values in the incidents log, 240 and 585 in the open and closed problem logs, respectively), does not map uniquely to the *Owner Country* field. We surmised that some names might be used by multiple people in different countries. We concatenated the owner countries with first names to create a new field, *Concatenated Country / Name* with 1,688 distinct entities for the incidents log, and 254 and 631 for the open and closed problems logs, respectively. We did not take into account the possibility that multiple distinct entities within the same country might be using the same name—this was not possible without additional information such as employee identification numbers.

Calculation of Step and Case Duration

Each event is associated with a single time stamp (the instant at which a status change occurs), so we determined time elapsed by calculating the difference between status changes. Under this convention, the final status in each case (usually *Completed*) is considered to conclude instantaneously.

Separation of Sub-statuses by Resource Input

We segmented portions of case duration associated with productive time and unproductive time (not requiring input by human resources) by status and sub-status for analysis of operational performance and productivity. We considered the sub-status *In Progress* as productive time spent working on the case. We considered *Queued–Awaiting Assignment* and *Accepted–Assigned* as unproductive case time with no active involvement by IT resources. The different *Wait* statuses (e.g. *Wait* or *Wait–Implementation*) had insufficient supporting data to determine whether this was

unproductive time or time when associated organizations were providing assistance to Volvo IT. We included *Wait* statuses in our calculations of total time but in our analysis it was treated as “Other” time, neither productive nor unproductive. The status *Completed-Resolved* is time after a solution has been delivered and is thus neither productive nor unproductive time for Volvo IT.

Product Groupings

Ideally, we would approach this type of data by consolidating products into broader categories wherever possible to make the data set more manageable. However, this is not possible due to the lack of identifying information about products, such as functionality and design. This also prevents us from drawing conclusions based on the nature of work being performed.

Linking Problems to Incidents

One of the most interesting pieces of analysis we would like to conduct is determining the causal factors for behind the elevation of incidents to problems. However, this type of analysis requires a larger incident data set that encompasses cases closed before the month of May.

Extraction of Service Line

The service line information is embedded within the Support Team (ST) designations themselves, as most values in the *Involved ST* field (for example, N52 2nd) contain both a support team number and a service line designator. We extracted these values and assumed entries without an explicit service line designator belong to Service Line 1, the common name for the Service and Expert Help Desks.

While most support teams are confined to handling events within a single service line, some STs do span across several lines, particularly within the incident management organization (**Table 2**):

Service Line Involvement	# of Support Teams (Incidents)	# of Support Teams (Open Problems)	# of Support Teams (Closed Problems)
1 Only	201	23	45
2 Only	255	117	186
3 Only	91	47	88
1 and 2	34	0	0
2 and 2.5	1	0	0
2 and 3	16	0	2
Grand Total	598	187	321

Table 2: *Service line* involvement for support teams in each of the three data sets

Hierarchy Assumptions

Our understanding of the organization hierarchy stems from the description provided [3] and our analysis of the data set. We suggest the structure for mapping support teams to their respective Organization or Function: Organization or Function → Support Team → Resource. This reflects the given definition of Organization as the business area of the user reporting the problem and Function as the groupings of IT divisions. Our analysis also supports this understanding given that the Organizations and Functions do not map one-to-one. Furthermore, Support Teams do not map one-to-one with Organization or Function and Resources do not map one-to-one with Support Teams.

4 Results

4.1 Process Conformance between Organizational Lines

We took a process centric approach to evaluating the conformance between organizational lines was process centric. This approach required a description of a standard process and descriptions of the standard process flows for Volvo IT incident and problem management were not included in the documentation. Through our initial analysis we discovered a standard process for both incident and problem management. This enabled us to evaluate the conformance of Organization Line (Org Line) A2 and C to the standard process and to each other.

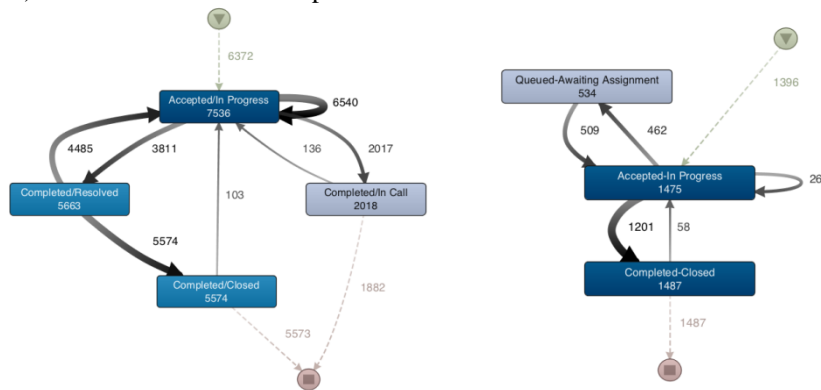


Figure 1: Highly simplified process maps of Incidents and Closed Problems. Left: Disco-generated process map of all incidents; Right: Disco-generated process map of all closed problems. Threshold settings, both process maps: Activities 20%, Paths 40%

Determining Standard Case Flow

We discovered the standard incident and problem management processes by leveraging Disco's built-in process map generator with activities and paths thresholds set to 20% and 40%, respectively. This gave us a highly simplified depiction of the

All Incidents	Start	Accepted-In Progress	Accepted-Wait	Accepted-Wait User	Accepted-Assigned	Queued-Awaiting Assignment	Completed-Resolved	Completed-Closed	End
	Start	-	6,372	-	-	-	1,156	-	-
Accepted-In Progress	-	6,540	1,092	2,303	1,723	-	3,811	-	-
Accepted-Wait	-	-	64	215	72	243	-	-	-
Accepted-Wait User	-	-	-	233	142	482	1,377	-	-
Accepted-Assigned	-	1,821	-	-	98	-	-	-	-
Queued-Awaiting Assignment	-	4,485	-	-	-	663	-	-	-
Completed-In Call	-	136	-	-	-	-	-	-	1,882
Completed-Resolved	-	-	-	-	-	88	-	5,574	-
Completed-Closed	-	103	-	-	-	-	-	-	5,573
End	-	-	-	-	-	-	-	-	-

Figure 2: Org Line C handles the vast majority of cases terminating in *Completed-In Call*. First second Org Line C's adjacency matrix and second all incidents' adjacency matrix. The number denotes the number of cases for which the sub status in the column followed, at any point, the sub status in the column.

When we compared these adjacency matrices to the simplified process map, we notice that of the 1,882 cases concluding with *Completed-In Call*, 1,800 (95.6%) are cases that involve Org Line C alone (Figure 2). The *Completed-In Call* designation is used whenever a service request is completed during a call to the help desk (Service Line 1). This suggests that Org Line C is the primary line responsible for Help Desk cases as it resolves roughly 96% of all cases that finish via *Completed-In Call*. In contrast, Org Line A2, and all the other Org Lines, handle very few help desk cases.

Ratio of Case Frequency Incidents involving Only Org Line C / Org Line A2	Start	Accepted-In Progress	Accepted-Wait	Accepted-Wait User	Accepted-Assigned	Queued-Awaiting Assignment	Completed-Resolved	Completed-Closed	End
	Start	-	1.0	-	-	*	1.1	-	-
Accepted-In Progress	-	1.0	1.1	1.2	1.8	1.4	-	-	-
Accepted-Wait	-	-	*	*	*	*	-	-	-
Accepted-Wait User	-	-	-	*	-	1.1	-	-	-
Accepted-Assigned	-	1.8	-	-	*	-	-	-	-
Queued-Awaiting Assignment	-	1.3	-	-	-	0.7	-	-	-
Completed-Resolved	-	-	-	-	-	*	1.0	-	-
Completed-Closed	-	-	-	-	-	*	-	1.0	1.0
End	-	-	-	-	-	-	-	-	-

Figure 3: Org Line C and Org Line A2 use most sub statuses with roughly the same frequency. This matrix shows the ratio of the normalized case frequencies. Ratios greater than one denote steps more frequently used by Org Line C, and less than one denotes steps more frequently used by Org Line A. Asterisks denote transitions traversed by less than 5% of cases for both Org Line A2 and C.

To further assess the difference between Org Lines A2 and C, we eliminated cases ending in *Completed/In Call*. We generated process maps with the same threshold levels as before (Activities 55%, Paths 35%). Since Org Line C handles many more cases than Org Line A2 we normalized the number of transitions by the total number of cases. We compared the two Org Lines by simply taking a ratio of the normalized number transitions, see Figure 3 above.

With few exceptions the ratios are near 1 which indicates that there is little other deviation between Org Lines A2 and C. We noted that Org Line C tends to utilize the sub-status *Queued-Awaiting Assignment* roughly 30% more frequently than Org Line A2 and the sub-status *Accepted-Assigned* nearly twice as often (Figure 3). We also note that both Org Line A2 and C conform the standard process from *Accepted-In Progress* → *Completed-Resolved* → *Completed-Closed*.

Process Conformity with Respect to Closed Problems

There are far fewer closed problems than incidents, so we were able to examine the process maps at a much greater level of detail. For these particular process maps, we set the activities and paths thresholds at 100% and 90%, respectively. We used the same normalization and ratio analysis we used on incidents excluding the cases ending in *Completed-In Call*.

Ratio of Case Frequency Problems involving Only Org Line C / Org Line A2	Start	Accepted-In Progress	Accepted-Wait	Accepted-Assigned	Queued-Awaiting Assignment	Completee-Cancelled	Unmatched	Completed-Closed	End
Start	-	1.0	*	*	*	-	-	-	-
Accepted-In Progress	-	0.9	0.9	0.8	2.4	*	*	1.1	-
Accepted-Wait	-	0.4	*	-	*	-	-	0.8	-
Accepted-Assigned	-	0.9	-	*	-	-	-	*	-
Queued-Awaiting Assignment	-	2.1	-	-	*	-	-	-	-
Completee-Cancelled	-	*	-	-	-	-	-	-	-
Unmatched	-	-	-	-	*	-	-	*	-
Completed-Closed	-	-	-	-	-	-	-	-	1.0
End	-	-	-	-	-	-	-	-	-

Figure 4: Org Line C has twice the proportion of problems that have sub status *Queued-Awaiting Assignment*. This matrix shows the ratio of the normalized case frequencies. Ratios greater than one denote steps more frequently used by Org Line C, and ratios less than one denote steps more frequently used by Org Line A. Asterisks denote transitions traversed by less than 5% of cases for both Org Line A2 and C.

Discussion of Process Conformance Analysis

We established standard process flows for both incidents and problems. We demonstrated that the primary difference between Org Line A2 and C is that Org Line C handles the vast majority of Completed–In Call incident cases. Excluding these cases the processes are roughly equivalent, with few exceptions. When handling incidents Org Line C has a rate of Accepted–Assigned 80% higher than Org Line A2. When handling problems a case handled by Org Line C is twice as likely to use the sub status Queued–Awaiting Assignment while a case handled by Org Line A2 is twice as likely to use the sub status Accepted–Assigned (Figure 4). Beyond these differences, Org Lines A2 and C seem to follow the standard incident and problem management process.

4.2 Push-to-Front

Our Understanding of Push-to-Front

Push-to-front (PTF) behavior is defined as incident reaching resolved by first-line personnel (hereafter referred to as *Service Line 1*) without involvement from higher-line support teams (*Service Lines 2 and 3*). Push-to-front resolution is preferred in modern IT incident management as it minimizes interruption of the duties normally performed by Service Lines 2 and 3, which typically do not include product support.

We analyzed the push-to-front issue though segmenting cases by initial org line, function, product, and country of origin. This strategy allowed us to identify org lines that handle primarily PTF (e.g. Org Line C) and elevated start cases (e.g. Org Line A2), as well as recognize the fact that a majority of functions are centered on elevated start incidents. We also identified products that are particularly prone to push-to-front resolution, and those that might benefit from reassignment to other case types.

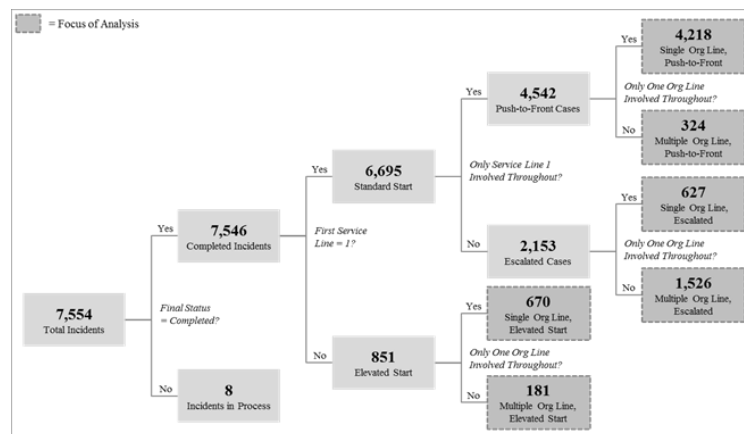


Figure 5: Distribution of all 7,554 incidents into one of six resolution types. Incidents classified according to org line composition and push-to-front behavior

With this in mind, we classified each completed incident into one of six resolution types, as depicted in **Figure 5**.

Push-to-Front Behavior by Initial Org Line

We segmented the completed incidents by their initial org lines, in order to more fully understand the role of these lines in handling cases of a particular resolution type (**Table 3**). The types of cases assigned to org lines varies significantly. With respect to Org Lines C and A2 (the lines to which ~86% of cases are initially assigned), the former primarily handles single org line, push-to-front cases (3,777 total incidents, 65.7% of line total) whereas the latter focuses mainly on single org line cases originating at Service Lines 2 or 3 (“escalated cases”; 362 total incidents, 48.7% of line total).

Org Line B appears to exhibit “C-like” behavior in that 168 of its 290 initially assigned cases (57.9%) are eventually resolved in a single org line, push-to-front fashion. One interesting observation is that these B-assigned incidents originate primarily from countries outside the European Union and North American regions, namely Australia, Brazil, China, India, Malaysia and Russia. Perhaps this line serves as an auxiliary support unit equipped to handle routine PTF incidents, so as not to overwhelm other lines handling more complicated calls from higher volume regions.

First Org Line	# of Completed Cases	Push to Front, Single Org Line	Push-to-Front, Multiple Org Lines	Escalated, Single Org Line	Escalated, Multiple Org Line	Elevated Start, Single Org Line	Elevated Start, Multiple Org Lines
C	5,746	66%	3%	9%	20%	2%	0%
A2	744	16%	7%	10%	10%	49%	9%
Other	419	27%	17%	0%	55%	0%	1%
B	290	58%	3%	11%	17%	11%	0%
G4	157	0%	0%	0%	0%	73%	27%
V2	69	45%	0%	41%	7%	7%	0%
G2	37	14%	22%	0%	65%	0%	0%
V5	21	0%	0%	0%	14%	33%	52%
G1	18	0%	0%	0%	0%	50%	50%
F	14	0%	0%	0%	36%	7%	57%
V11	13	0%	23%	0%	54%	0%	23%
H	7	71%	0%	0%	14%	0%	14%
Misc.	11	0%	0%	0%	27%	0%	73%
Total	7546	56%	4%	8%	20%	9%	2%

Table 3: Distribution of completed incidents by first Org Line and resolution type. Leading resolution types for each Org Line are highlighted in **bold**.

While most escalated cases are handled as part of the portfolio of incidents managed by Org Line C (1,129 total incidents, 19.6% of total), a number of org lines appear to also specialize in the handling these incidents, notably G2, V2 and the undesignated org line “other”, which itself could be instances of org lines listed elsewhere.

Push-to-Front Behavior by Support Team Function / Division (ST Function Div)

In a manner similar to that used for org line analysis, we segmented all 7,546 completed incidents by the identity of their initial functions (**Table 4**).

In contrast to the varied specializations exhibited by the different org lines, most of the functions (17 of 21) appear to specialize in single org line, elevated start incidents. Only two, E_5 and V3_2 have strong specialization tendencies toward push-to-front cases, while a third, A2_1, has a fairly even distribution between these types of incidents and escalation (specifically multiple org line) cases.

First ST Function Div	# of Completed Cases	Push to Front, Single Org Line	Push-to-Front, Multiple Org Lines	Escalated, Single Org Line	Escalated, Multiple Org Line	Elevated Start, Single Org Line	Elevated Start, Multiple Org Lines
V3_2	4,802	70.8%	2.7%	9.2%	17.1%	0.0%	0.1%
A2_1	986	38.6%	11.9%	9.4%	33.4%	5.1%	1.6%
_	836	19.7%	9.0%	3.5%	30.5%	27.2%	10.2%
E_5	421	60.8%	0.0%	12.1%	25.2%	1.9%	0.0%
A2_4	159	0.0%	0.0%	8.2%	1.3%	78.0%	12.6%
D_1	89	0.0%	0.0%	0.0%	5.6%	91.0%	3.4%
A2_2	73	1.4%	0.0%	0.0%	0.0%	74.0%	24.7%
E_6	50	4.0%	0.0%	0.0%	10.0%	46.0%	40.0%
A2_3	48	0.0%	0.0%	0.0%	0.0%	87.5%	12.5%
A2_5	25	40.0%	4.0%	0.0%	0.0%	44.0%	12.0%
E_10	21	4.8%	0.0%	0.0%	4.8%	76.2%	14.3%
C_6	13	0.0%	0.0%	0.0%	0.0%	100.0%	0.0%
E_1	5	0.0%	0.0%	20.0%	0.0%	60.0%	20.0%
E_8	5	0.0%	0.0%	0.0%	0.0%	80.0%	20.0%
Misc.	13	0.3%	0.0%	0.0%	0.5%	92.3%	7.7%
Total	7546	55.9%	4.3%	8.3%	20.2%	8.9%	2.4%

Table 4: Distribution of completed incidents by first ST Function Div and resolution type. Leading resolution types for each function are highlighted in **bold**. Misc includes all other STs not otherwise included in the table.

Product	# of Completed Cases
---------	----------------------

Product	# of Push-to-Front Cases	% of Completed Cases Exhibiting
---------	--------------------------	---------------------------------

566	158
832	39
369	30
505	20
420	19
522	15
732	15
533	14
794	14
53	13

Table 5: Top 10 products exhibiting 100% PTF behavior

		PTF Behavior
424	684	77.6%
660	442	91.3%
383	193	94.1%
253	172	76.4%
566	158	100.0%
494	142	76.3%
13	107	81.7%
321	94	87.0%
267	79	66.4%
453	77	83.7%

Table 6: Products with the highest number of push-to-front incidents

Push-to-Front Behavior by Product

Since many products are represented by only a single incident we sought to simplify our analysis by setting a minimal case threshold while still representing most incidents, resulting in 226 products comprising 6,724 of the original 7,546 completed cases (89.1% coverage). Under this threshold, a number of products exhibit strong PTF behavior, both by the proportion and absolute number of PTF cases handled (Tables 5 and 6).

On the opposite end of the spectrum, there are a number of products that begin their lives at Service Line 1 but are eventually escalated to the higher service lines prior to completion. For some of these particularly high-volume products (**Table 7**), perhaps a re-designation as “elevated start” may prove beneficial in terms of time saved and a decrease in overall complexity for these cases.

Product	# of Completed Cases
542	75
604	36
295	32
337	28
54	27
818	27
308	20
488	18
631	18
591	17

Table 7: Top 10 products exhibiting 0% PTF behavior

As highlighted in the VINST user manual, “recording a solution also makes it possible for you to resolve similar SR’s without doing extensive research” [5]. However, “Solutions are objects in the database that are separate from Service Requests” and require users to link associate solutions to service requests. To promote push-to-front behavior and decrease total work we could require or incentivize the addition and association of solutions to all incidents.

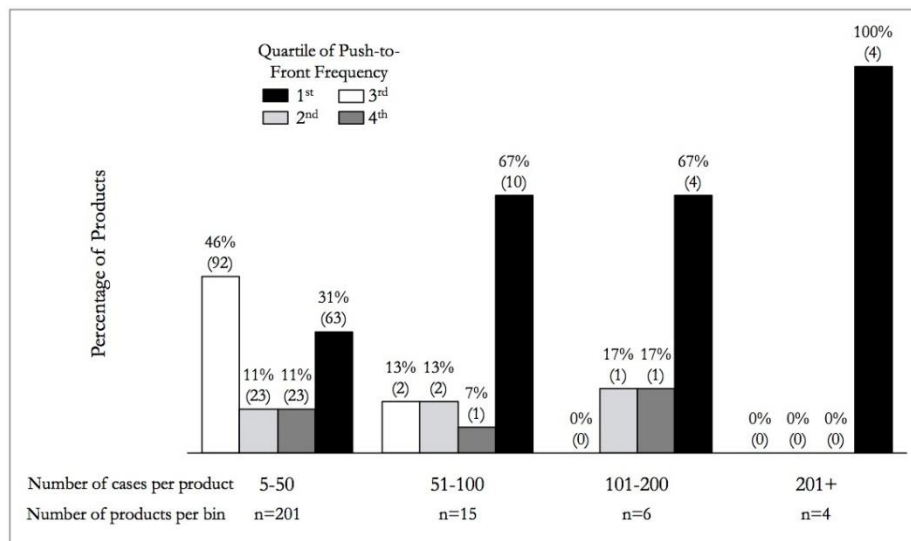


Figure 6: Correlation between push-to-front behavior and number of completed incidents, grouped by product. Products are first segmented into bins by number of cases (5-50, 51-100, 101-200, and greater than 201) and then divided into quartiles by push-to-front frequency (0-25%, 25-50%, 50-75%, and 75-100% of cases exhibiting push-to-front). The percentage of products per quartile and number of products per quartile are indicated above each bar.

Finally, we focused on the PTF behavior of individual products, and learned that some are certainly more prone to PTF resolution (**Figure 6**) than others (**Table 7**). Additionally, we evaluated the behavior of those incidents that are not push-to-front in nature (that is, involve some escalation to Service Lines 2 or 3 during their lifetime) to see if any opportunities exist for streamlining the handling of these cases. For example:

- Are there specific products for which a large percentage of reported cases begin at Service Line 1, but are eventually escalated to Service Lines 2 or 3?
- Are there cases that spend a very short time at Service Line 1 before being escalated?

To answer these questions, we categorized the products according to the percentage of cases that are eventually escalated from Service Line 1, and the average elapsed time

(in minutes) from the first recorded event to the time of escalation. We then filtered out those products containing fewer than five total incidents (see above) and isolated those with $\geq 75\%$ escalated cases (out of a total of 10 or more completed incidents) and < 20 minutes spent at Service Line 1 prior to first escalation.

Product	# of Completed Cases	% of Cases Escalated to Service Lines 2 or 3	Average Time (minutes) Spent at Service Line 1 Prior to First Escalation	Standard Deviation, Time Spent at Service Line 1 Prior to First Escalation
488	18	100.0%	1.3	5.6×10^{-4}
238	11	81.8%	4.0	2.5×10^{-3}
431	10	90.0%	4.2	1.6×10^{-3}
726	14	100.0%	10.2	5.5×10^{-3}
655	10	90.0%	11.1	9.6×10^{-3}
542	75	98.7%	11.5	1.6×10^{-2}
305	29	93.1%	17.6	1.7×10^{-2}

Table 8: Potential candidates for re-designation as elevated start products

To this end, we identified seven products (**Table 8**) that might benefit from re-designation to higher service lines upon submission to the incident management system, bypassing the initial handling at Service Line 1 and possibly streamlining their resolution as “elevated start” products.

Push-to-Front Behavior by Country of Origin

Finally, we examined the push-to-front behavior of cases by country of origin, in order to evaluate whether specific countries may be responsible for a disproportionate number of cases from a certain resolution type. As shown in **Figure 7** below, three countries, Poland, USA and Brazil, are strong originators ($\geq 65\%$ of cases belong to a single resolution category) of push-to-front incidents, while a fourth, Canada, produces almost exclusively elevated start cases:

While $> 98\%$ of incidents from the Netherlands are escalations, all but two of these belong to a single product, 542, which suggests that this behavior is due to the product itself rather than on practices being followed by the country’s reporting staff. However, in Canada, a location from which a much larger number of products (27 in total) is reported, 65 of 66 ($> 98\%$) are elevated start cases, thus suggesting the opposite. As the ambiguous nature of the source data makes it difficult to make hypotheses about the products themselves, additional information is necessary in order to make definitive conclusions about the behavior of countries toward the incidents they handle.

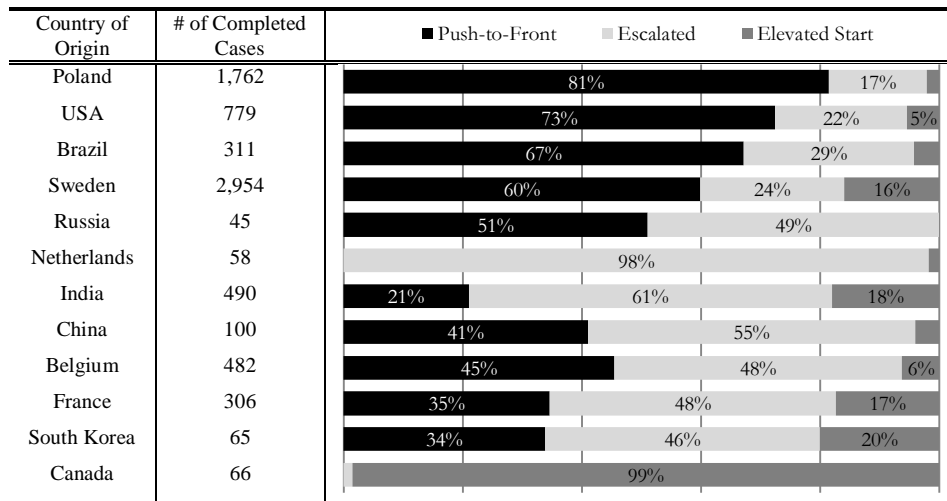


Figure 7: Push-to-front behavior of completed cases by country of origin. Not included in this table are low-volume countries Malaysia (26 cases), Australia (25), UK (22), Japan (15), Thailand (6), Chile (3), Peru (2) and Turkey (2). Also not included: 26 cases with no country of origin attached. Total coverage, excluding these entities: 98.3%

Conclusions and Potential Opportunities for Improvement

Our analysis identified the specialization tendencies of the org lines initially assigned to cases in the incident management system, which proved to be widely distributed among push-to-front, escalated, and elevated start cases. Most importantly, we identified a distinct difference between the cases initially assigned to Org Lines A2 (mainly elevated start) and C (push-to-front), the two lines to which ~86% of cases are assigned upon submission to VINST. We also learned that while a majority of the functions (17 of 21) specialize in elevated start cases, only three of the remaining four have a tendency to handle push-to-front incidents.

With respect to countries of origin, we learned that the nature of cases handled vary widely between the three resolution types, but given the state of the source data, were unable to conclude whether these observations were due to the product and problem profiles of the reporting locations or the reporting tendencies of the countries themselves. A more detailed description of the resolution process would allow us to measure process conformity of cases belonging to similar or identical products, and evaluating the overall process by which all incidents are handled.

4.3 Ping Pong Behavior

Our Understanding of Ping Pong Behavior

Many cases are handled with the involvement of a single support team while others require the involvement of additional support teams to reach a satisfactory resolution. We analyzed the occurrence of ping pongs between support teams, which we defined as any time a support team works on a specific case more than once following a transfer between support teams. This definition accounts for direct ping pongs ($A \rightarrow B \rightarrow A$) and indirect ping pong cycles ($A \rightarrow B \rightarrow C \rightarrow A$). The documentation for BPIC 2013 describes ping pongs as an “unwanted situation” so we leveraged the provided data to assess the business impact of ping pongs [3].

We first evaluated the impact of ping pongs on case completion time, and then determined which support teams are most responsible for ping pongs. We identified the products with the highest ping pong rate and the impact that a targeted initiative aimed at reducing ping pongs could have on the work time of support teams at Volvo IT. We recognize that certain processes may require support teams to transfer a case back and forth for various legitimate reasons; however, we are unable to distinguish definitively between legitimate ping pongs that follow process design and illegitimate ping pongs from this data. Nonetheless, we point to concentrations of activity.

The Impact of Ping Pongs on Completion Time

To determine the effect of ping pongs on incident duration, we compared the mean durations of incidents with ping pongs to those without. To ensure that our calculations were independent of variation stemming from product differences, we analyzed incidents (5,893 incidents, 78% of total) concerning the 205 products that had incidents with and without ping pongs. These results were split into deciles as shown in Figure 8.

Cases exhibiting ping pong were on average 2.3-fold longer than those without ping pongs while holding product constant and excluding the top and bottom deciles (Figure 8). Incident-weighted mean case durations were 201.0 and 465.6 hours, for those without and with ping pongs, respectively. Both the first and tenth decile warrant further analysis as ping pong cases in the top decile are associated with case durations orders of magnitude above the other nine deciles.




Figure 8: The mean duration of incidents with ping pongs is longer than incidents without ping pongs for c.80% of products. Product deciles by fold-change in mean case duration for when cases exhibit at least one ping pong (n=205). The top decile is plotted on right-side axis. The dotted line indicates equivalency of mean case duration with and without ping pong. *Other* includes all status and sub-status designations not specifically listed.

We also determined the average portion of time incidents spend in steps of each sub-status. We considered *In Progress* to be the status that identifies the actual work effort by the Volvo IT support teams and *Queued* and *Assigned* as indications of unproductive process time that increase the total resolution time. The distribution of time spent in steps of these sub-statuses varies across deciles.

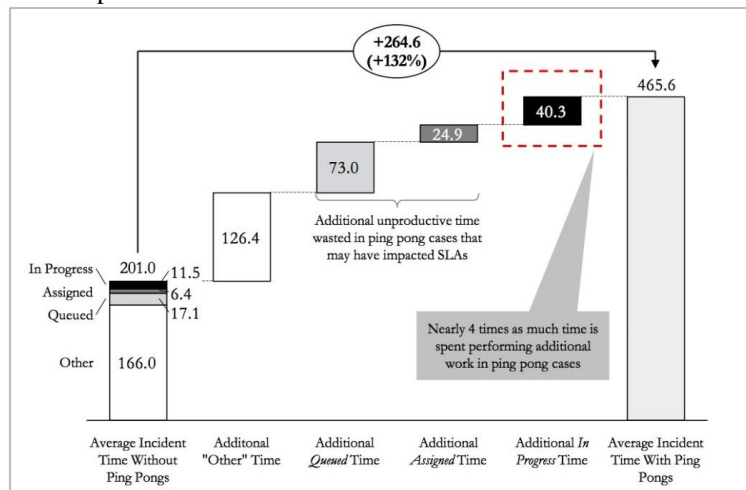


Figure 9: Ping pongs increase the average case time to roughly 2 times while increasing work time and unproductive wait time to roughly 5 times. Source of change analysis between average incidents without ping pongs and average incidents with a ping pong. Average times are in hours and the additional times are the difference between the average time for each respective status in the average incident without a ping pong and the average time for this status designation in an average ping pong case.

The increase in the time spent in these steps has two separate but related impacts: 1) increased *In Progress* time adds work for support teams, and 2) increased *Queued, Assigned* and *Other* time increases the total incident time affecting customer and potentially SLAs. The total impact of the additional 40.3 *In Progress* hours for 952 cases is 38,366 hours of support team work. A significant portion of this work is likely unnecessary rework that burdens support teams and reduces productivity. Likewise, customers waited an additional 93,201 hours for cases to complete. This likely impacts service level agreements (SLAs), and we would be able to more fully assess this impact with descriptions of the SLAs.

Ping Pong Activity by Support Team

We computed both the absolute number of ping pongs and their frequency attributable to support teams, Org Lines, and functional divisions. The ping pong frequency was defined as the ratio of ping pong events to the total number of transfers between STs.

	# of Transfers	# of Ping Pongs	Ping Pong Frequency	Function	Org Line
Service Line 1 (total = 231)	11,113	2,107	0.19		
D4	546	336	0.62	A2_1	A2
G97	1,233	236	0.19	V3_2	C
D5	408	210	0.51	A2_1	C
D8	402	170	0.42	A2_1	A2
D2	346	162	0.47	A2_1	C
D7	205	106	0.52	A2_1	A2
D1	181	103	0.57	A2_1	B
G96	1,686	69	0.04	V3_2	C
D6	124	61	0.49	A2_1	C
G92	226	58	0.26	E_5	C
S49	148	58	0.39	V3_2	C
Line 1 Subtotal (5% of total)	5,505 (50%)	1,569 (75%)	0.29		
Service Line 2 (total = 306)	4,500	974	0.22		
V37	235	159	0.68	-	V7n
N18	50	32	0.64	A2_5	A2
N14	67	31	0.46	A2_1	A2
Line 2 Subtotal (1% of total)	352 (8%)	222 (23%)	0.63		
Service Line 3 (total = 107)	751	92	0.12		
G42	45	14	0.31	A2_1	A2
G107	10	5	0.50	A2_4	A2
Line 3 Subtotal (2% of total)	55 (7%)	19 (21%)	0.35		
Grand Total	16,384	3,173	0.19		

Table 9: A small number of support teams in each service line are responsible for the majority of ping pong events and often also have a high frequency of ping pong. The top 5% of support teams by total ping pongs for Service Line 1, the three support teams from the top 5% of Service Line 2 (by total ping pongs) that have the highest ping pong frequency, and the top two support teams by total ping pong from Service Line 3. *Excluded:* Service Line 2.5

An examination of total ping pongs by support team identified a strong concentration of ping pongs (Table 9). 73% of ping pongs are attributable to a 5% of STs. Segmenting further, the top 1% of STs (D4, G97, D5, D8, D2, V72 2nd, and D7) are responsible for 43% of the total ping pong events (1379 ping pongs, 21% of all transfers). These support teams generally also have high ping pong frequencies (Table 9). The six of the seven STs in the top 1% belong in Service Line 1, which is responsible for 68% of all ping pongs. However, the trend of a few service teams accounting for the majority of ping pong events holds true in each service line.

This demonstrates that a small number of support teams are responsible for a majority of the ping pong events across several dimensions. These teams also are strongly associated with functional division A2_1 and Org Line A2. Future analysis should focus on the roles, functions, organization, and connections between these groups, determine the root cause of ping pongs, and optimize incident management processes.

Ping Pong Activity by Product

We also examined ping pongs by product to determine which products are most affected by ping pong. Of the 701 named products, 255 products had at least one incident with a ping pong event, which we segmented by ping pong frequency into deciles. Since we observed that the ping pong frequencies of the top decile are dramatically higher (1.6-fold from the 9th to 10th), we focused our analysis on the constituent 21 products (in order of frequency from 0.80-0.57, Products 510, 799, 736, 141, 775, 303, 158, 727, 137, 157, 398, 97, 542, 789, 776, 159, and 558).

Short-term change efforts should be focused on the teams that contribute the most ping pongs to this top decile of products. We compare the ping pong frequency of STs for a product with the product's average ping pong frequency. A subset of STs repeatedly exhibited above average frequencies—D4, D5, D2, D7, D6, D8, D1, G97 and G57 2nd all exhibited ping pong frequencies above the product average for at least two products (10, 8, 6, 5, 3, 3, 2, 2, and 2 products, respectively). This is perhaps not surprising given that many teams work on the same products.

This analysis arrives at a similar conclusion to the STs analysis of ping pongs by identifying the same teams as most responsible for ping pong behavior. The key next step would be to engage these teams to understand their contribution to incident resolution such that these ping pongs can be categorized into essential and nonessential (and therefore noncompliant). Furthermore, additional information about the products—especially product relationships and hierarchies—would also be beneficial in determining whether specific products or product groups have qualities that predispose them to ping pong behavior.

Discussion of Ping Pong Analysis

Through our analysis we have identified that a small number of support teams disproportionately contribute to the total number of ping pong events and the 38,366 hours of *In Progress* work time. From this information we believe that Volvo IT could conservatively reduce total support team work time by 10,600 hours (28%).

We propose that a targeted initiative focusing on the eight support teams from Service Line 1, the three teams from Service Line 2 and the two support teams in Service Line 3 listed in Table 9 which are most responsible for ping pongs. Together these teams have a ping pong frequency of 0.65. By reducing this rate by one third to 0.4 (still double the average frequency across all support teams) Volvo IT would eliminate 700 ping pongs a month (given that this data set is representative of a typical month), which, with an average of 2.66 ping pongs per incident with ping pongs, equates to 236 incidents worth of ping pongs. This reduction would save 40 hours of *In Progress* time per incident for a total of 10,600 hours for the month. Furthermore, this analysis does not address the top decile of products, shown in Figure 9, where other initiatives could yield substantial additional savings.

A properly designed initiative would determine the root cause of ping pongs through further analysis of roles, functions, organization, and connections between the support teams, and org lines. Additionally, as certain org lines and functions are disproportionately represented by the support teams identified in our analysis such an initiative must be aimed at the appropriate organizational level in order to affect its intended outcome.

4.4 Wait-User Analysis

Our Understanding of the Wait-User Issue

According to the BPIC 2013 documentation, *Wait-User* is a sub status (under the *Accepted* category) used by Action Owners to manually “stop the clock” on a particular case in order to decrease the total turnaround time for completion of a task. While there are certainly some legitimate uses for this sub status (such as waiting for information or action from a user), some owners are suspected of blatantly misusing it as a means of improving their own performance metrics in the incident and problem management systems. Cases that include the use of *Wait-User* status have a 7-fold longer case duration with 20% of the additional case time due to *Wait-User* (Figure 10).

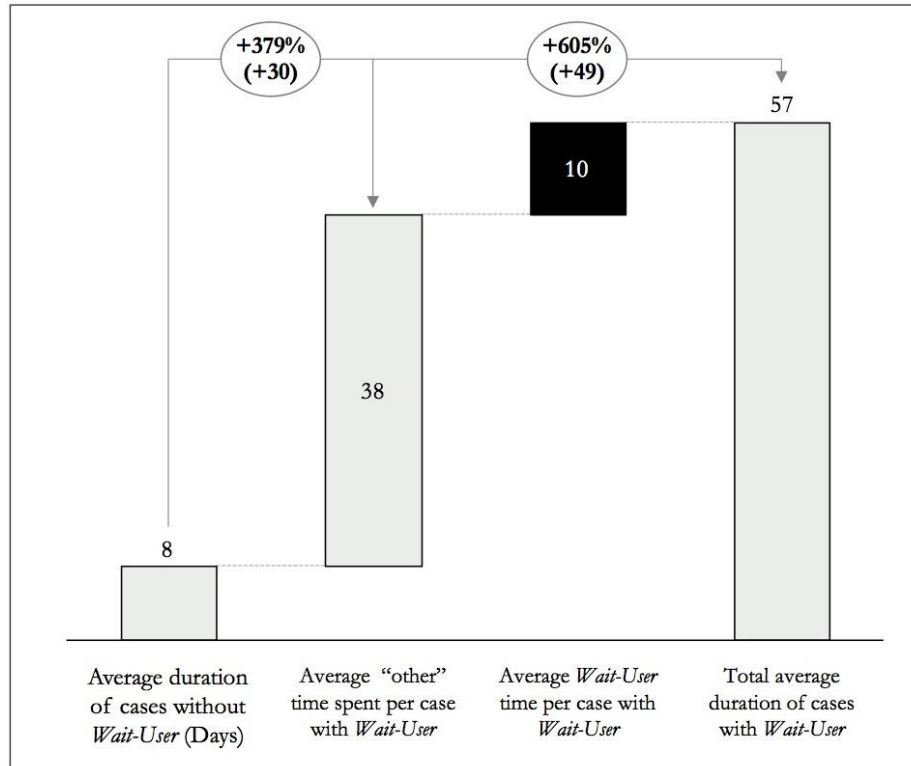


Figure 10: *Wait-User* cases are 7-fold longer in average duration with 17% of the increase due to *Wait-User* time. Source of change analysis between average incidents without *Wait-User* and average incidents with *Wait-User*. Average times are in days and the additional times are the difference between the average time non-*Wait-User* cases.

In order to understand the use of *Wait-User* time by individual action owners, we investigated its usage across the various STs, Org Lines, Functions and owners by country. For this analysis we subset the data and only examined cases that included the *Wait-User* sub status. We calculated the average usage of the *Wait-User* sub status per case as well as the total time that a case spent in this sub status.

Wait-User Distribution by Support Team

To evaluate ST performance with regards to the use of *Wait-User*, we analyzed the average frequency of *Wait-User* usage and the duration of *Wait-User* by team. Teams that consistently use *Wait-User* and do so for longer periods of time across many products have a disproportionately high impact upon

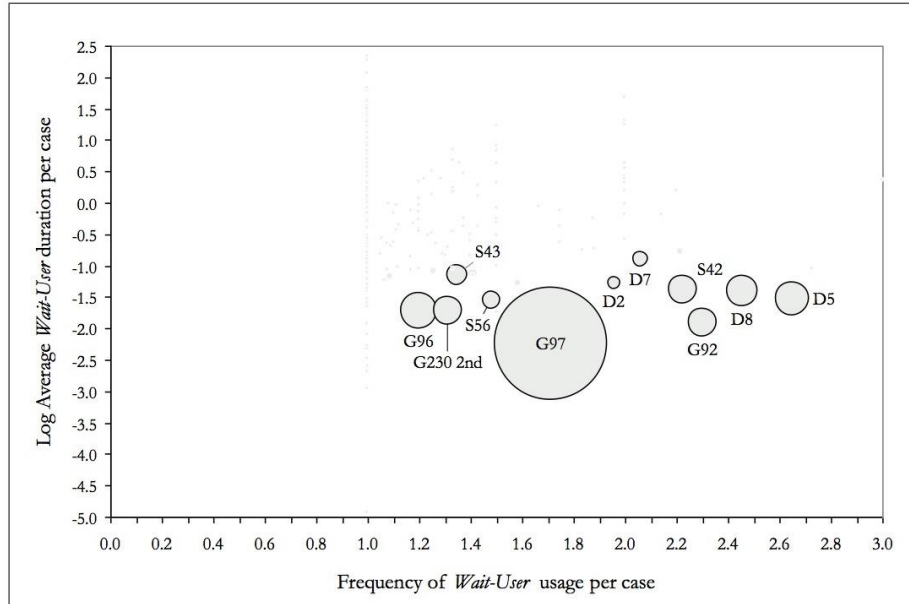


Figure 11. Average *Wait-User* frequency and duration for STs that work on the most products. *Wait-User* duration is plotted on a log scale and the size of the circle is proportionate to the number of products on which the ST has worked.

From this analysis we observe that STs in Service Line 1 frequently use *Wait-User* and for longer durations than their peers. While it is difficult to conclude that these instances are definitely related to abuses of the *Wait-User* functionality, any meaningful examination of this phenomenon would surely begin with an investigation of these five support teams.

We conducted an analysis of STs by product to identify teams that disproportionately impact average *Wait-User* time at the product level. We identified 8 teams that have an average *Wait-User* time more than 7 days above the average for that product on multiple products (**Table 10**).

ST	Count of Products Where ST's Mean <i>Wait-User</i> Duration Exceeds Overall Average by ≥ 7 days (Percentage of all such cases)
N49	3 (5.08%)
S24	3 (5.08%)
G96	2 (3.4%)
G97	2 (3.4%)
M10	2 (3.4%)

N25	2 (3.4%)
S41	2 (3.4%)
S55	2 (3.4%)

Table 10: STs underperforming other teams' *Wait-User* averages by ≥ 7 days

The combination of the above analyses demonstrate that two STs, G96 and G97 (highlighted in Figure 11 and **Table 10**), not only have the highest number of deviations from the average *Wait-User* time when all products are considered, but also exhibit *Wait-User* durations that exceed their peers by more than a week. However, in order to draw conclusions regarding legitimate use of Wait User time, we would require more specific information about the process requirements, ST capabilities, and the formal guidelines regarding proper usage of the *Wait-User* sub status.

Wait-User Distribution by Org Line

Analysis of *Wait-User* activity for the different Org Lines exhibited more consistency in the frequency of *Wait-User* usage and in the mean duration of total *Wait-User* time per case. The Org Lines that had the highest frequency of *Wait-User* usage had a lower mean duration of *Wait-User* usage although the variance of the distribution of these frequencies was low. Org lines C and A2 have the highest usage of *Wait-User* time but a lower mean duration compared to V2 and B, both of which have fewer occurrences.

Org Line C has higher overall frequency of *Wait-User* usage, whereas Org Line A2 has a longer average duration per instance of *Wait-User*. This may be due to the initial assignment of more complex, elevated start cases to Org Line A2 as described in our push-to-front analysis.

Wait-User Distribution by Support Team Function / Division

We analyzed *Wait-User* activity by Function assess usage at this organizational level. These groups also exhibited consistency in the frequency of *Wait-User* usage and in the mean duration of total *Wait-User* time per case similar to what we observed in our analysis of Org Lines. Only A2_3 had a substantially longer average *Wait-User* time per case. Interestingly, the three functional divisions with the highest frequency of *Wait-User* usage are generally the initial group assigned to routine cases exhibiting push-to-front type behavior.

Wait-User Distribution by Location

To understand the usage of *Wait-User* by action owners by location we examined the frequency of *Wait-User* usage and in the mean duration of total *Wait-User* time per case. As with our analyses of Org Lines and functions, it appears that owners who use the *Wait-User* status the most often are not the ones who use it for the longest

duration. While high usage is consistent across Sweden, the owners using the option the most in Sweden do not have the highest average *Wait-User* time per case.

Discussion of Wait-User Analysis

The use of *Wait-User* option correlates well with significantly longer case durations, however the use of the option stops the clock from ticking as far as turnaround time for a case is concerned. This allows for the option to be open to abuse at the user level. Our analysis of *Wait-User* activity demonstrated clear outliers at the ST, Org Line and Owner-Country level. ST's G97, G96, G230 2nd, D7 and D8 use the *Wait-User* option the most across the most products and use it for durations longer than other ST's working on the same products. Org Lines that use the option the most exhibit lower durations of *Wait-User* usage. Org Line C has experiences higher usage than Org Line A2, which may be indicative of initial assignment of more complex elevated start cases to A2.

5 Conclusions

Our analysis has identified several performance improvement opportunities in the IT incident management process for Volvo, Belgium. Such improvements would increase resource efficiency and decrease customer impact through reduction of case resolution times. We believe the potential performance gains are substantial and warrant further investigation and analysis to develop specific action plan to realize such gains.

We discovered the standard process maps for both incidents and problem across Org Lines A2, Org Line C, and the other Org Lines. From these maps, we identified 1,800 Help Desk calls belonging exclusively to Org Line C, and separated these from tickets not Completed-In Call. Excluding calls, we found similar processes performed in Org Line C and A2, which indicates some degree of standardization between these Org Lines. This suggests the presence of a business reference model—likely captured in a common language (e.g. business process model and notation)—that provides a focal point for the process modifications this study recommends. In the absence of a reference model, the processes we discovered could be easily translated into formats easily leveraged by internal process owners.

We were able to identify process inefficiencies related to push-to-front, *Wait-User* abuse, and ping pong. In general, substantial opportunity for improvement is concentrated in a small number of support teams. The identified noncompliance of small group can be further investigated and gains realized in the near-term.

We observed a strong correlation between push-to-front activity and incident frequency for products—that is products with more incidents are less likely to be escalated. This tendency to keep more frequent products at Service Line 1 indicates a potential to better leverage knowledge management practices for less frequent

products. A learning curve appears to exist for product-specific solutions to be delivered by dedicated first line resources—the more incidents related to a product, the more likely the product is to remain at the 1st line. We hypothesize that there is a knowledge sharing mechanism (formal or informal) in place for the most commonly occurring incidents. VINST documentation indicates recording solutions to incidents and problems is an optional step [5] in the process. Creating incentives to capture such solutions formally is likely to help improve push-to-front behavior and reduce ping pong behavior.

The prevalence of ping pong behavior appears to cause significantly decreased resource efficiency and increased resolution times. We recommend further investigation into specific STs to segment ping pongs required for incident or problem resolution. Given the large size and concentration of this opportunity, there is clear opportunity to reduce ping pongs by identifying and addressing team-specific root causes.

Our analysis of *Wait-User* activity demonstrated clear outliers at the ST, Org Line and Owner-Country level. Among ST's G97, G96, G230 2nd, D7 and D8 stood out as teams using the Wait-User option the most. These teams stood out in terms of having selected Wait-User for the most products and for the longest durations among ST's working on the same products. While outlier behavior is apparent from our investigations, any inquiry into the legitimate or illegitimate use of *Wait-User* behavior requires a better understanding of the roles of the different ST's, Org Lines, Owners and the different Products lines.

We identified 255 cases (c. 3%) with total duration over 50 days. In general, we excluded from analyses when considering total duration. These cases could represent data quality or logging issues because the majority of time appears to be spent Queued—Awaiting Assignment (i.e. resources resolved a case, but failed to record the resolution). However, if engagement with involved Action Owners revealed these cases are reflective of ongoing issues, understanding the details of these cases would provide substantial value for improving SLA performance (and accuracy) and reducing customer impacts.

Overall, given the requirement of making data publicly available, Volvo shared very limited data about the incidents and problems. We understand the limitation placed on sharing any additional data given risks with sharing confidential information. However, if similar analyses are carried out with access to relevant details about incidents, products, problems, resources and the organization structure, and performance expectations; one could build highly actionable process / operations change recommendations to drive meaningful performance improvement.

Our conviction in the power of process mining of event logs, in combination with details about the work items and the overall operational set up to yield powerful insights has grown further with our participation in this year's BPIC challenge. We thank Volvo for making the data available and the organizers of this competition for allowing us to participate.

Acknowledgements: We thank Lalit Wangikar, Nick Hartman and Eric Chung for their discussion and helpful suggestions.

Contributions: AB performed Push to Front analysis, SA performed *Wait-User* analysis, TM performed ping pong analysis, and AA performed process conformity analysis. AB, SA, AA, TM and MLR interpreted data and composed manuscript.

References

1. Ferreira, D. and Mira da Silva, M., Using process mining for ITIL assessment: a case study with incident management. In: Proceedings of the 13th Annual UKAIS Conference, Bournemouth University.
2. Jochen De Weerd, Seppe K. L. M. vanden Broucke, Jan Vanthienen, Bart Baesens: Leveraging process discovery with trace clustering and text mining for intelligent analysis of incident management processes. IEEE Congress on Evolutionary Computation 2012: 1-8
3. VINST data set and description. Jun 14, 2012.
http://www.win.tue.nl/bpi2013/lib/exe/fetch.php?media=vinst_data_set.pdf. Accessed: Jul 10, 2013
4. "BPIC Challenge 2013." Flux Capacitor. Fluxicon, Jun 1, 2013. Accessed: Jul 11, 2013.
5. VINST User Manual. Jun 14, 2013.
http://www.win.tue.nl/bpi2013/lib/exe/fetch.php?media=vinst_manual.pdf. Accessed: Jul 11, 2013.