

Rabobank: A Process Mining Case Study

BPI Challenge 2014 Report

Gabriele Cacciola¹, Raffaele Conforti², and Hoang Nguyen²

¹ Università della Calabria, Italy
gabrielecacciola90@gmail.com

² Queensland University of Technology, Australia
raffaele.conforti@qut.edu.au
huanghuy.nguyen@student.qut.edu.au

Abstract. The increased interest in the area of process improvement persuaded Rabobank Group ICT in examining its own Change-process in order to improve its competitiveness. The group is looking for answers about the effectiveness of changes applied as part of this process, with particular interest toward the presence of predictive patterns and their parameters. We conducted an analysis of the log using well established process mining techniques (i.e. Fuzzy Miner). The results of the analysis conducted on the log of the process show that a visible impact is missing.

1 Introduction

Companies want to be competitive on the market, and efficiency is often one of the requirements for competitiveness. In order to achieve efficiency it is important to avoid bottlenecks, idle times, and delays in business processes. The identification of these problems, which may affect the efficiency of a process, can be carried out through a process mining [2] exercise. Similarly to other companies, Rabobank Group ICT is interested in knowing if its processes are efficient and effective.

In order to align IT services with the needs of business, Rabobank Group ICT implemented the ITIL-processes. Among these processes, we can find the Change-process which is used for implementing changes in IT services. Rabobank Group ICT is particularly interested in knowing if changes, realized a part of its Change-process, impact the workload at the Service Desk and/or IT Operations.

Specifically, they are interested in finding answers to the following questions³:

1. **Identification of Impact-patterns:** We expect there to be a correlation between the implementation of a change and the workload in the Service Desk (SD) and/or IT Operations (ITO), i.e. increased/decreased volume of Closed Interactions and/or increased/decreased volume of Closed Incidents. Rabobank Group ICT is interested in identifying any patterns that may be visible in the log for various service components to which a configuration

³ www.win.tue.nl/bpi/2014/challenge

item is related, in order to predict the workload at the SD and/or ITO after future changes.

2. **Parameters for every Impact-pattern:** In order to be able to use the results of prior changes to predict the workload for the Service Desk directly after the implementation of future changes, we are interested in the following parameters for every impact-pattern investigated in question 1:
 - What is the average period to return to a steady state?
 - What is the average increase/decrease of Closed Interactions once a new steady state is reached?
3. **Change in Average Steps to Resolution:** Since project managers are expected to deliver the same or better service levels after each change implementation, Rabobank Group ICT is looking for confirmation that this challenge is indeed being met for all or many Service Components.
4. **Creativity challenge:** Finally, we challenge the creative minds, to surprise Rabobank Group ICT with new insights on the provided data to help change implementation teams to continuously improve their Standard Operation Procedures.

The analysis conducted in order to answer these four questions produces several interesting insights. In particular, an unexpected trend in the process was observed and in contrast with what expected from Rabobank Group ICT the actuation of changes does not reduce the chance of new disruptions leaving the number of interactions almost unchanged.

The rest of the paper is structured as follows. Section 2 discusses the logs and how they were extracted from the four datasets. Section 3 presents the techniques we used in order to answer the four questions presented here. Section 4 discusses the findings of our analysis and Section 5 concludes the paper.

2 Data Extraction and Filtering

Before presenting the log provided by Rabobank Group ICT and describing the techniques used for extracting and filtering the logs, we will briefly describe of the process. Whenever a customer encounters a disruption with an ICT-service, he will have an interaction with the Service Desk. In the large number of cases, the disruption can be solved by the Service Desk clerk during the first call, if a first call resolution is not possible an incident will be issued and the intervention of an expert is required. If an issue recurs several times, a problem and improvement analysis is planned in order to understand the real cause and prevent future service disruption. After the completion of the analysis, a change is implemented on the configuration item, responsible of the issue.

Rabobank Group ICT distributed the log of the Change-process as four datasets, covering the period between the 01/10/2013 to the 31/03/2014. The datasets contain the Interaction records (counting 147004 events), the Incident records (counting 46606 events), Incident Activity records (counting 466639 events), and the Change records (counting 30236 events).

Using these datasets, four logs were created: IT Operations, Service Desk, Change, and Incident. Table 1 shows the number of cases and events of each log created. The IT Operations log and the Service Desk log contain records of interactions, incidents (only for IT Operations), and changes occurred as part of the Change-process. Finally, the Change log contains events related with the execution of changes, while the Incident log contains events related with the management of incidents.

Log	Number of Cases	Number of Events
IT Operations	8811	1166938
Service Desk	13691	177240
Change	18000	124130
Incident	46606	466665

Table 1. Number of cases and events for each log.

Among the four logs we mentioned, two of them, the Change log and the Incident log, were created by Fluxicon⁴ and we will not discuss them.

The remaining two logs have been obtained as follow. The Service Desk log is based on an SQL query, which merges events of the datasets containing changes and interactions. Table 2 shows the query used for the generation of the log.

```
create table logsd
(
  select CIName as caseid, CIName as CIName, CI_Type, CI_Subtype,
  Change_ID as ID, tab, Actual_Start as STime, Actual_End as ETime,
  Service_Comp_WBS as serviceComponent, idchange as id_table
  from change
)
union all
(
  select CIName as caseid, CIName as CIName, CI_Type, CI_Subtype,
  interaction_ID as ID, tab, open_time as STime, close_time as ETime,
  Service_Comp_WBS as serviceComponent, idinteraction as id_table
  from interaction
)
```

Table 2. Query for the generation of the Service Desk log.

The IT Operations log contains the chain of events that from an interaction, led to the creation of an incident and finally to a change. Based on our understanding of the process, when an interaction occurs it is associated with a configuration item (see CIName in the interaction table) which is expected be the cause of the disruption. The occurrence of an interaction may generate an incident, which will be linked to the configuration item generating the interaction (see CIName_Affected field). As consequence of several incidents of

⁴ www.fluxicon.com

the same type a change may be required. This change will be performed on a configuration item, which may not be the configuration item generating the interaction. Finally, the configuration item affected by the change will be linked with the incident (see `CI_Name_CausedBy` field).

The description of the synergy between interaction, incidents and changes shows that an incident may be related with two different configuration items, one generating the interaction and another been affected by a change. Using this insight we decided to generate the IT Operations log based on the `CI_Name` of a change. Figure 1 shows the chain of `CI_Names` used, for each `CI_Name` we identified related changes, we then selected all incidents caused by it, and finally we selected all interactions generated by the `CI_Name_Affected` of each incident.



Fig. 1. configuration item name chain: Change → Incident → Interaction

During the extraction of the logs we had to make several choices, one of them was deciding what timestamp to use for our activities. The four datasets provided offer a total of sixteen different fields that could be used as possible timestamps. Table 3 shows the list of all possible timestamps. We decided to use the following timestamps: i) *Open Time (First Touch)* and *Close Time* as start and end time for an interaction, ii) *Open Time* and *Resolved Time* as start and end time for an incident, and iii) *Actual Start* and *Actual End* as start and end time for a change.

Dataset	Field	Start Time	End Time
Interaction	Open Time (First Touch)	✓	
	Close Time		✓
Incident	Open Time	✓	
	Reopen Time		
	Resolved Time		✓
	Closed Time		
IncidentActivities	DateStamp		
Change	Planned Start		
	Planned End		
	Scheduled Downtime Start		
	Scheduled Downtime End		
	Actual Start	✓	
	Actual End		✓
	Requested End Date		
	Change record Open Time		
	Change record Close Time		

Table 3. Selected timestamps among all possible timestamps.

Finally, we applied a set of filters on the logs in order to remove eventual noise. Two types of filters were applied on the Service Desk log and on the IT Operations log. Firstly, we removed all cases having events with unreliable timestamps (i.e. events with missing timestamps, or events with timestamps in the future), then we removed all cases not starting with either a change or an interaction event. In case of the IT Operations log we applied an additional filter, and we removed all cases not containing at least one event for each of the three types of activity (i.e. incident, interaction, and change), while for the Service Desk log we applied the same filter requiring occurrence of events belonging to the interaction and change activities.

The Change process log, on the other hand, was filtered removing all cases not containing at least one event for activities Actual Start and Actual End, this to maintains only cases where a change was performed and completed.

3 Data Manipulation

In the Introduction we presented the four questions Rabobank Group ICT is interested in. In order to answer those questions, we first analyzed the log looking for changes in the number of interactions and incidents resulting from the application of a change. This was achieved using a simple algorithm that, each time an event happens, plots the current number of interactions and incidents along with the number of changes occurred so far. An example of plot can be seen in Figure 2, this type of graph will allow an easy identification of trends in the behavior of interactions and incidents in response of a change.




Fig. 2. Fluctuation in the number of active interactions and incidents (blue and red), in response of a change (green).

In order to address the second question we are interested in measuring the average period required to return to a steady state, and the difference in the number of closed interactions/incidents between a change and another. In the context of this work, a change reaches a steady state as soon as the number of closed interactions/incidents is equal to the average number of closed interactions/incidents minus one standard deviation. The average number of interactions/incidents and its standard deviation are computed across the time between

a change and the following one. Figure 3, shows how the steady state is computed. In particular, each change (i.e. $c1$, $c2$, $c3$, and $c4$) reaches a steady state when the number of interactions (black line) reaches the green line. The time needed to reach this state (i.e. $TimeSS$) is represented by the teal line, while δ represents the difference in the number of closed interactions/incidents between two changes.

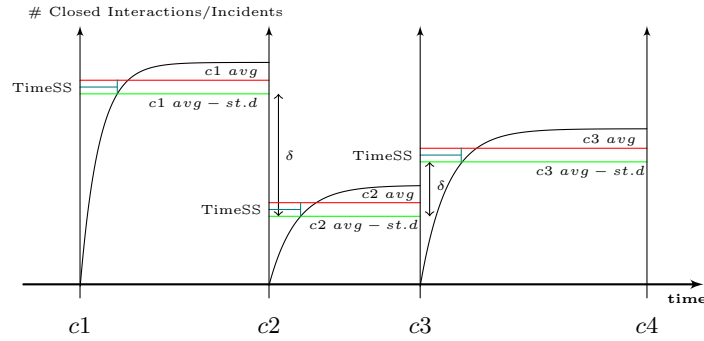


Fig. 3. Example of steady state computation.

Answering question three requires to know if for a given service component, an increasing number of changes will gradually lead to a lower number of interactions between changes. This information can be obtained using the same algorithm that we used for question one, but in this case we consider all changes occurred to configuration items been part of the same service component.

Finally, in order to answer question four, we used the process mining tool Disco⁵. We subdivided the Incident process log in several logs, where each log contains all incidents related with configuration items of the same type. Each of the logs so created was filtered removing all infrequent activities (i.e. appearing with a frequency below 1%), and we mined the fuzzy model [1] limiting the number of paths to 35% of the total paths.

4 Results

We are now in the conditions to proceed with the discussion over the findings of our analysis. Our findings show that the Change-process presents a behavior that contrasts with what is expected from Rabobank Group ICT. In the following we analyses each question one by one illustrating our findings.

⁵ www.fluxicon.com/disco/

4.1 Identification of Impact-patterns

The identification of impact-patterns requires plotting the number of active interactions/incidents over the cumulative number of changes using the technique illustrated in Section 3. We performed our analysis over both the IT Operations log and the Service Desk log.

IT Operations Log: Considering the large number of configuration items affected by changes, we decided to group them based on the configuration item type. Figure 4 shows the variation in the number of active interactions and incidents for configuration items of type software. The plot shows the absence of a clear pattern in response to the actuation of a change. We would expect drops in the blue (number of active interactions) and red (number of active incidents) lines every time an increase in the green (cumulative number of changes) occurs, while this pattern clearly is missing.




Fig. 4. Fluctuation in the number of active interactions and incidents (blue and red), in response to a change (green) for configuration items of type software for the IT Operations log.

On the other side, the plot shows an noticeable increase in the number of active interactions in response to a sequence of changes. This sequence of changes was probably the cause of a disruption, since it required an additional sequence of changes in order to mitigate the situation and reduce the number of active interaction to the norm. A similar behavior can be observed for all remaining configuration item types, for example Figure 5 shows the variation in the number of active interactions and incidents for configuration items of type computer, for the plots of all remaining configuration item types we refer to the Appendix.

The plots also show another interesting trend. In the log we could observe that for certain configuration item types changes are not executed in response to an interaction or a change, since they are executed when the number of active interactions and incidents is equal to zero. In particular, this configuration




Fig. 5. Fluctuation in the number of active interactions and incidents (blue and red), in response to a change (green) for configuration items of type computer for the IT Operations log.

item types are: hardware (see Figure 6), applicationComponent, database, and networkComponent. We refer to the Appendix for the plots of the remaining configuration item types.

Service Desk Log: The Service Desk log presents a trend similar to the IT Operations log. In this case, we can observe (see Figure 7) that no clear pattern can be identified as consequence of a change.

In this case, we can see that changes are not executed in response to interactions, see for example the graph in Figure 8 showing the fluctuation in the number of active interactions and incidents for configuration items of type hardware.

4.2 Parameters for every Impact-pattern

Rabobank Group ICT is interested in knowing the average period needed to return to a steady state after a change, and the average increase/decrease number of closed interactions once a new steady state is reached. The answer to this questions can be found using the second technique presented in Section 3.

Average period to steady state

IT Operations Log: The previous analysis showed the absence of patterns in the variation of the number of interactions and incidents in response to changes. Despite the absence of patterns we performed an analysis to figured out if after a change a steady state is actually reached and eventually after how long. This analysis required, as the previous analysis, to group configuration items based on their type.




Fig. 6. Fluctuation in the number of active interactions and incidents (blue and red), in response to a change (green) for configuration items of type hardware for the IT Operations log.




Fig. 7. Fluctuation in the number of active interactions and incidents (blue and red), in response to a change (green) for configuration items of type application for the Service Desk log.




Fig. 8. Fluctuation in the number of active interactions and incidents (blue and red), in response to a change (green) for configuration items of type hardware for the Service Desk log.

The result of our analysis showed that after a change a steady state is never reached. Based on the description of Rabobank Group ICT we were expecting the cumulative number of closed interactions to grow following a trend similar to the cumulative distribution function of the exponential distribution. The cumulative number of closed interactions/incidents instead grows linearly as can be see in Figure 9. Figure 9 shows the cumulative number of closed interactions (blue line) and incidents (red line) between changes for configuration items of type computer, we can observe that the average number of closed interactions (green line) is reached half way between two changes.




Fig. 9. Cumulative number of closed interactions (blue line) and incidents (red line) between changes for configuration items of type computer for the IT Operations log.

A similar trend can be observed for each configuration item type, Figure 10 shows The cumulative number of closed interactions/incidents for configuration

items of type subapplication, we remand to the Appendix for the remaining configuration item types.




Fig. 10. Cumulative number of closed interactions (blue line) and incidents (red line) between changes for configuration items of type subapplication for the IT Operations log.

Service Desk Log: The Service Desk log, as for question one, presents a behavior similar to the IT Operations log. In this case we cannot identify a steady state. As we can observe from the plot in Figure 11 the number of closed interactions grow linearly. The same behavior can be observed for the remaining configuration item types, we refer to the Appendix for the related plots.

Increase/Decrease in the number of closed interactions: In order to answer the second point of question two, the presence of a steady state is required. As we have already observed, since the number of closed interactions and incidents follows a linear growth, no steady state is reached between changes. Based on this, it is not possible to investigate the increase/decrease in the number of closed interactions when a steady state is reached.

4.3 Change in Average Steps to Resolution

In the third question, Rabobank Group ICT is interested in knowing if after each change the service level of Service Components is the same or better. The analysis conducted so far showed how changes on a configuration item do not affect the number of disruptions caused by it. On the basis of this finding we expect that the service level of Service Components will follow a similar distribution, since a service component is composed of multiple configuration items.

In order to verify if service components behave similarly to configuration items we used the third algorithm described in Section 3. The result of our




Fig. 11. Cumulative number of closed interactions (blue line) and incidents (red line) between changes for configuration items of type application for the Service Desk log.

analysis produced over 300 different diagrams. Visualizing all of them would have been quite difficult, for this reason we decided for each service component to compute how many times the average in the number of closed interactions between two changes increases or decreases. Figure 12 shows for each service component the number of times the average increased and decreased between two consecutive changes. We can observe that service components tend to behave as configuration items not following a predefined behavior as expected from Rabobank Group ICT. In particular, we can observe that quite often the average between changes increases and decreases independently from the number of changes applied.




Fig. 12. Increase (blue line), decrease (red line), and difference (green line) in number of closed interactions between changes for each service component for the IT Operations log.

A similar trend can also be observed in case of the Service Desk log, as showed in Figure 13.




Fig. 13. Increase (blue line), decrease (red line), and difference (green line) in number of closed interactions between changes for each service component for the IT Operations log.

4.4 Creativity challenge

The analysis conducted so far showed that the Change-process does not reflect the expectations of Rabobank Group ICT. Based on this, we decided to provide some additional insights about the process. In particular, we grouped the processes by configuration item type and for each of them, we discovered the process executed when an incident is issued. These models are then used to perform a delta analysis to measure and identify differences between processes. Finally, we analyzed the resource utilization and we investigated the process in order to identify possible bottlenecks.

Delta Analysis The discovery of model was performed using the technique describe in Section 3. The incident activity log was filtered, maintaining only the activities occurring with a frequency equal to or greater then 1%, while we keep only the most frequent paths (appearing in 65% of the cases). After mining the process of each configuration item type, we noticed that there are processes having activities that others do not have, see for example the process shows in Figure 14 and in Figure 15.

As shown in the pictures Figure 14 and Figure 15, activities such as Pending Vendor, Communication with customer and External Vendor Assignment exists only for particular configuration item types. This can be explained if we consider that for the handling of different configuration item types different approaches are used. For example, in case of display, hardware, and network component, external entities are contacted to obtain spare parts from suppliers.

Bottlenecks and Resource Utilization In an organization is very important to identify activities that are resource drainers. Analyzing the performance of the process, we noticed that for certain activities the idle time before their execution is quite long. In particular, this situation occurs with three different configuration item types. Figure 16 shows the incident activity process in case of configuration item type Application. In the process we can observe that activity Operator Update, executed in 3725 case, has an average idle time of 9 hours, while the median duration of the entire process is equal to 3 days.




Fig. 14. Process model mined for the configuration item type Display




Fig. 15. Process model mined for the configuration item type Display

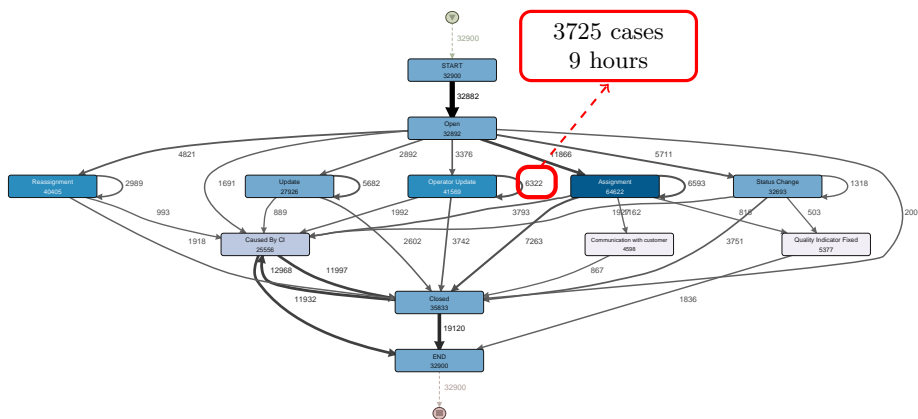


Fig. 16. Bottleneck for configuration item type application

Figure 17 shows another example, in this case we have the process for configuration item type computer. We can observe that in 596 cases after the execution of activity Status Change the process has to wait on average 22.6 hours over a median duration of the entire process of 54.7 hours.

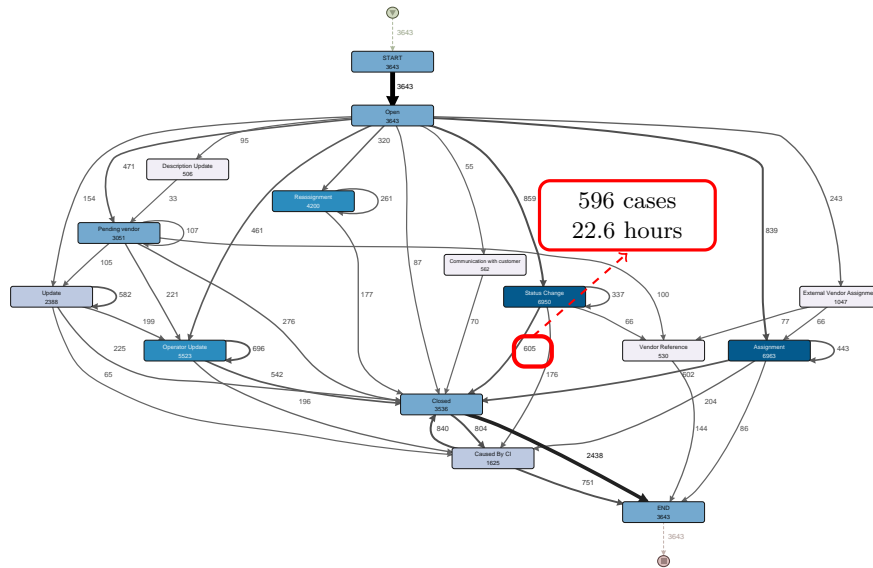


Fig. 17. Bottleneck for configuration item type computer

As part of our resource utilization analysis we observed that multiple teams are involved in the same incident handling process. In the process, we noticed that on average for each incident there are 2-3 teams (2.62 teams) involved with the handling of the process. However, there are many incidents where a large number of teams is involved, leading us to question if this may affect the performance the incident handling process.

Two process metrics are used to evaluate the incident process: *Actual handle time* and *Meet SLA*, which are all based on the *Handle Time* field of an incident. The handle time, based on the provided documentation, is the "time registered to resolve the incident, based on a SLA clock mechanism in HP Service Manager 9". Actual handle time is computed as the duration between *open time* and *resolved time* of an incident. We then compared the actual handle time with the handle time to determine if a case meets the SLA or not.

The graph in Figure 18 plots the actual handle time and SLA compliance over the number of teams involved (Meet SLA = Yes for red line and No for blue line). In general, we can observe that while the time needed to handle an incident increases with the number of teams involved, a drop in time (from 1200

to 750 hours) can be detected when ten teams are involved. Moreover, when the number of teams exceeds 6, the time needed to resolve the incident never meets the SLA (the registered handle time).




Fig. 18. Actual Handle time and MeetSLA over Number of teams involved.

We then investigated the actual handle time for different configuration types and sub-types. The circle in Figure 19 represents the length of actual handle time (the size of a circle is proportional to the handle time). It points out that some areas of configuration items take a considerably longer handle time than others, including: web-based application, MigratieDummy, Server-based Application, Standard Application. These software applications by nature are normally built on multiple layers of technologies which might require different teams with different skill sets and expertises to assist in troubleshooting incidents.

These findings persuaded us in investigating the impact of human resources on the incident handling process. In Disco, using resources as activities, we could draw a dependency graph among various teams as shown in Figure 20. We can observe that Team0008 assumes the role of coordinator during the handling of an incident, since there are many communication flows to and from this team. An additional analysis showed that Team0008 is involved in the handling of over 77% of incidents.

The graph in Figure 21 shows the daily number of incidents in which Team0008 is involved. It shows that Team0008 takes part in approximately 600 incidents per day (the drops on the graph represent weekends).

Figure 20 also reveals that the transfer of work between Team0008 and Team9999 takes an considerably large amount of time, as can be depict from the arc label denoting the median duration of the idle time. We then identified the couples of teams with the longest handover of work time, the results are




Fig. 19. Configuration type/subtype handle time matrix.




Fig. 20. Resources dependency graph.




Fig. 21. Workload for Team0008.

shown in Table 4. The table shows that Team9999, Team0007, and Team0008 are the teams with the highest handover of work time.

Sender	Receiver	Duration (hours)	Number of Occurrence	Number of Cases where it occurs
TEAM9999	TEAM0002	52.4 days	49	42
TEAM0007	TEAM0008	70.2 hours	3,172	436
TEAM0008	TEAM9999	64.9 hours	1,385	1,190
TEAM0007	TEAM0007	45.8 hours	11,727	2,593
TEAM9999	TEAM0008	42.8 hours	6,124	1,012
TEAM0086	TEAM0008	63.9 hours	211	80

Table 4. Longest handover of work time

Finally, our investigation shows that there are potentially three reasons why the communication between these teams is slow:

Waiting for customer update Notably, in most of the cases the high idle time taken for the handover of work between Team9999 and Team0002 is related to Update from customer activity.

Task processing Sender or Receiver may take time in reporting or analyzing. This is probably caused by the workload on either sender or receiver side, or by the complexity nature of the incident.

Re-open an incident An incident can be re-opened after it has been closed.

On the basis of the analysis conducted, we can observe that every time a handover of work takes longer Team0008 is involved. This identifies Team0008 as a possible bottleneck in the process, considering that this team has very high number of incidents per day (approximately 6 times larger than the others).

5 Conclusions

In conclusion, the analysis conducted shows that the execution of changes does not affect considerably the workload in the Service Desk and/or IT Operations, in contrast with what expected by Rabobank Group ICT. The analysis also shows that changes are not triggered in response to incidents/interactions, while no visible correlation can be detected between the triggering of a change and the number of active incidents/interactions. The absence of such correlation determines the absence of visible patterns for any service component which a configuration item is correlated to, making impossible to predict the workload at the SD and/or ITO after future changes.

The analysis also showed that between changes a steady state is never reached, since the number of close incidents/interactions grows linearly with the time. Moreover, the average number of closed incidents/interactions does not follow a specific trend, making impossible to confirm if the service level of service components improves after the implementations of changes.

Finally, we identify several bottlenecks in the process which are related with the execution of certain activities and/or specific teams.

References

1. Christian W. Günther and Wil M. P. van der Aalst. Fuzzy mining - adaptive process simplification based on multi-perspective metrics. In Gustavo Alonso, Peter Dadam, and Michael Rosemann, editors, *BPM*, volume 4714 of *Lecture Notes in Computer Science*, pages 328–343. Springer, 2007.
2. Wil M. P. van der Aalst. *Process Mining - Discovery, Conformance and Enhancement of Business Processes*. Springer, 2011.

A Appendix

The following figures illustrate the fluctuation in the number of active interactions and incidents, in response to a change for the IT Operations log.




Fig. 22. Fluctuation in the number of active interactions and incidents (blue and red), in response to a change (green) for configuration items of type application for the IT Operations log.




Fig. 23. Fluctuation in the number of active interactions and incidents (blue and red), in response to a change (green) for configuration items of type applicationComponent for the IT Operations log.




Fig. 24. Fluctuation in the number of active interactions and incidents (blue and red), in response to a change (green) for configuration items of type database for the IT Operations log.




Fig. 25. Fluctuation in the number of active interactions and incidents (blue and red), in response to a change (green) for configuration items of type networkComponents for the IT Operations log.




Fig. 26. Fluctuation in the number of active interactions and incidents (blue and red), in response to a change (green) for configuration items of type software for the IT Operations log.




Fig. 27. Fluctuation in the number of active interactions and incidents (blue and red), in response to a change (green) for configuration items of type storage for the IT Operations log.

The following figures illustrate the cumulative number of close interactions and incidents between changes for the IT Operations log.




Fig. 28. Fluctuation in the number of active interactions and incidents (blue and red), in response to a change (green) for configuration items of type applicationComponent for the Service Desk log.




Fig. 29. Fluctuation in the number of active interactions and incidents (blue and red), in response to a change (green) for configuration items of type computer for the Service Desk log.




Fig. 30. Fluctuation in the number of active interactions and incidents (blue and red), in response to a change (green) for configuration items of type database for the Service Desk log.




Fig. 31. Fluctuation in the number of active interactions and incidents (blue and red), in response to a change (green) for configuration items of type networkComponents for the Service Desk log.




Fig. 32. Fluctuation in the number of active interactions and incidents (blue and red), in response to a change (green) for configuration items of type officeelectronics for the Service Desk log.




Fig. 33. Fluctuation in the number of active interactions and incidents (blue and red), in response to a change (green) for configuration items of type software for the Service Desk log.




Fig. 34. Fluctuation in the number of active interactions and incidents (blue and red), in response to a change (green) for configuration items of type storage for the Service Desk log.




Fig. 35. Fluctuation in the number of active interactions and incidents (blue and red), in response to a change (green) for configuration items of type subapplication for the Service Desk log.

The following figures illustrate the fluctuation in the number of active interactions and incidents, in response to a change for the IT Operations log.




Fig. 36. Cumulative number of closed interactions (blue line) and incidents (red line) between changes for configuration items of type application for the IT Operations log.




Fig. 37. Cumulative number of closed interactions (blue line) and incidents (red line) between changes for configuration items of type database for the IT Operations log.




Fig. 38. Cumulative number of closed interactions (blue line) and incidents (red line) between changes for configuration items of type hardware for the IT Operations log.




Fig. 39. Cumulative number of closed interactions (blue line) and incidents (red line) between changes for configuration items of type networkComponent for the IT Operations log.




Fig. 40. Cumulative number of closed interactions (blue line) and incidents (red line) between changes for configuration items of type subApplication for the IT Operations log.




Fig. 41. Cumulative number of closed interactions (blue line) and incidents (red line) between changes for configuration items of type software for the IT Operations log.

The following figures illustrate the fluctuation in the number of active interactions and incidents, in response to a change for the Service Desk log.




Fig. 42. Cumulative number of closed interactions (blue line) and incidents (red line) between changes for configuration items of type computer for the Service Desk log.




Fig. 43. Cumulative number of closed interactions (blue line) and incidents (red line) between changes for configuration items of type hardware for the Service Desk log.




Fig. 44. Cumulative number of closed interactions (blue line) and incidents (red line) between changes for configuration items of type networkComponent for the Service Desk log.




Fig. 45. Cumulative number of closed interactions (blue line) and incidents (red line) between changes for configuration items of type subApplication for the Service Desk log.




Fig. 46. Cumulative number of closed interactions (blue line) and incidents (red line) between changes for configuration items of type software for the Service Desk log.