PROCESS MINING COLLECTIVE BEHAVIOR

Summary

Telefónica is a global telecommunications company, with more than 300 million customers in 21 countries. The management of their portfolio of digital services involves a variety of teams. They are complex organizations working around the clock, both globally and locally. These groups follow sets of rules and protocols, and their compliance can be analysed with data-driven techniques, including process mining.

Detecting operational drifts is a very challenging task. Slow changes in the informal culture of groups are not dramatic enough to produce a sharp impact on quality of service. They are not easy to detect, even for experienced analysts, because they do not change the overall process map. To make the problem harder, the analysis team has to survey dozens of different services, so it needs a reliable early indicator of anomaly.

In this article, we explain the strategy that we have followed to detect operational drift. 'Severity' is a primary parameter in incident management that indicates how severe an incident is. There are explicit rules for setting up the 'Severity' of an incident and for modifying it. Searching for abnormal percentages of 'Severity' change and unexpected concentrations in states of the process where they should be infrequent has proven to be an excellent proxy for detecting operational drifts and an informal cultural drift indicator.

We describe the search of the root cause in one major service. As a result, an unexpected operational practice was fixed.

- **Process mining case study in IT Service Management area**
- **A hidden cultural drift was detected and fixed**
- **Key success factor was to include 'Severity' changes into the process view**

Company

Telefónica I+D is the research and development company of the Telefónica Group. It was founded in 1988 and its mission is to contribute to the group's competitiveness and modernity through technological innovation. With this aim, the company applies new ideas, concepts and practices in addition to developing products and advanced services.

Our group is responsible for the definition and auditing of Operation and Deployment best practices within Telefónica I+D. We ensure the quality of the data generated by these activities and we define the analytics strategy of the area. We provide analytic capabilities to execute the defined strategy and advice to other units in the company.

Rationale

Conformance checking is one of the most important tasks in process analytics. Following the right protocol is important to achieve the standards of quality in any business, but sometimes it is not enough. When uncertainty is part of the recipe, there are unexpected scenarios that human teams solve adapting their informal work culture. Although they are not violations of the protocol, minor creative workarounds may lead to inefficient behaviour.

During the last year, we have been working on the automatic detection of early symptoms of operational drift across dozens of services. The analysis of time series has proven a good indicator of major changes like sharp increases in the traffic or number of customers. But there are also silent, minor modifications of daily behaviour that are not encoded but become part of the informal group culture. They are very difficult to detect, because rules are not explicitly broken and teams do not feel that these workarounds are potentially dangerous.

Analysing the patterns of change of one of the attributes of trouble tickets, the ‘Severity’ parameter, we have been able to find services with abnormal rates. We explain how this affected one of the services, and how it was fixed with the help of process mining.

Incident Life Cycle and Changes of ‘Severity’

The service that we use to explain our approach is an IoT managed connectivity service developed by Telefónica. By May 2017, more than 1 million of users of different sectors as automotive or e-commerce have been connected in Germany, Spain, Brazil, Argentina, Chile, and Mexico.

This service provides operational support seven days per week, 24 hours a day. Teams have a maximum of five hours to restore the service in case of a ‘Critical’ disruption, and twelve hours in case of a ‘Major’ degradation. The availability target is 99.5%.

An incident is defined as an unplanned interruption or reduction in the quality of a service. The incident management process ensures that normal service operation is restored as quickly as possible.

The canonical Incident life cycle is shown in Figure 1. Exceptionally, an incident can also be Cancelled or Delayed.

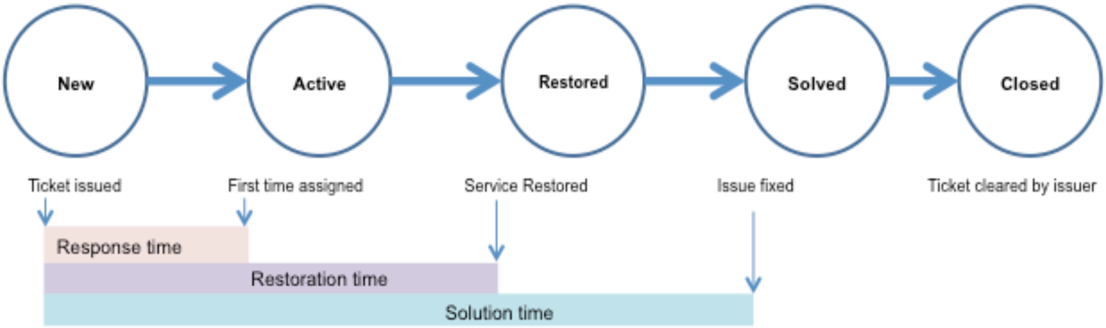

Figure 1: Typical life cycle of an incident.

As defined by ITIL, 'Severity' is the seriousness of an incident, and it is a function of the impact in the correct functioning of the service. We work with four levels of severity:

'**Critical**': A service disruption has happened and it requires immediate corrective action.

'**Major**': An incident is partially impairing service. It requires urgent correction.

'**Minor**': An incident impairs service but not seriously.

'**Slight**': The incident does not currently impair service, but the condition needs to be corrected before it becomes more severe.

Incident 'Severity' is not a constant attribute during the life of the incident. It is set up by the issuing group and may be modified by the operation teams following well-defined rules. Our hypothesis is that a high percentage of changes implies a weak initial categorization or hidden mismanagement.

Data

The data source is the in-house case management tool UDO. Information is retrieved using the native API. The system records all human interactions, so there is a wealth of information available.

Raw data can be downloaded and cleaned with R to select the relevant fields for process mining.

This pre-process is launched from a web interface built with the Shiny package (see Figure 2). The result is that ticketing data with the selected parameters are available for Disco processing.

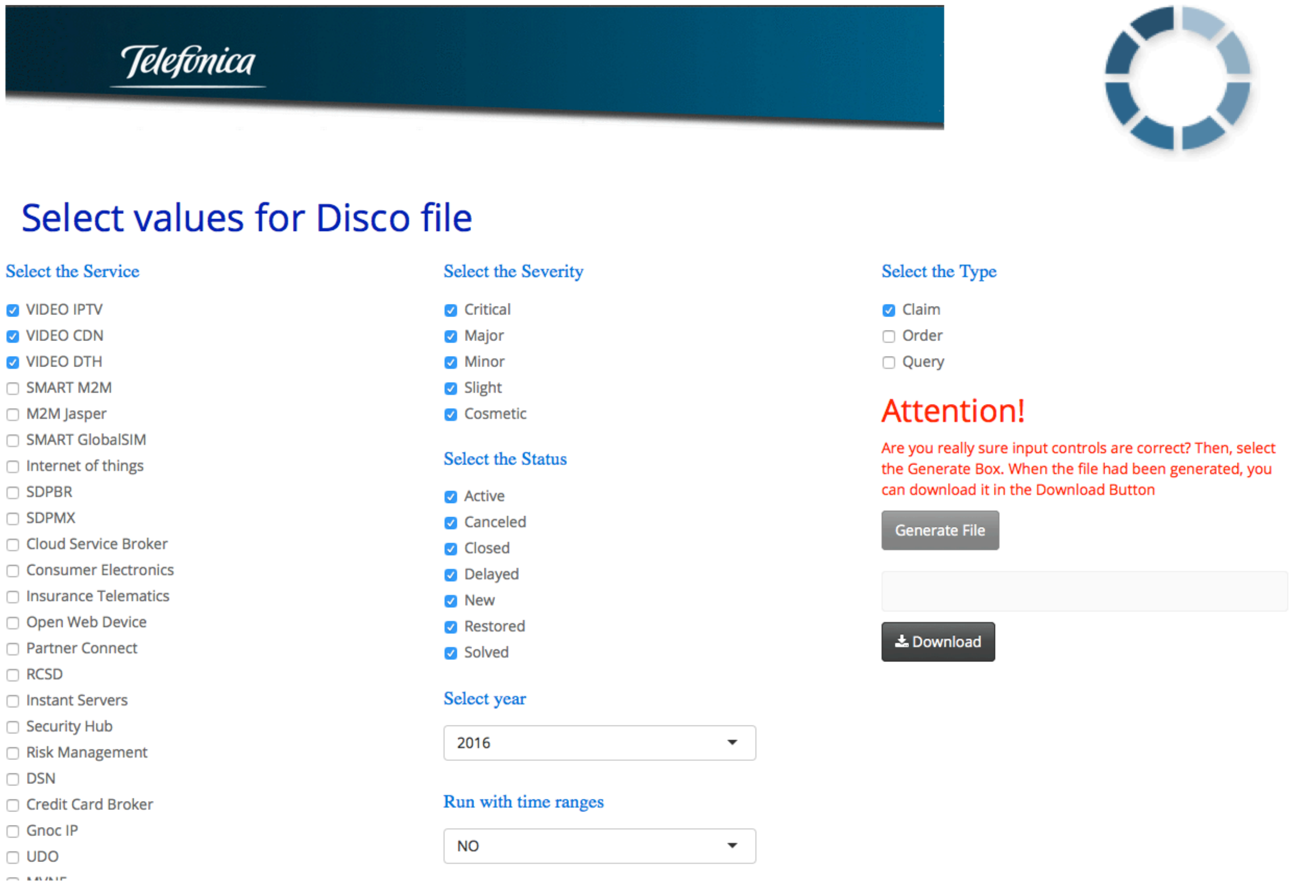

Telefonica

Select values for Disco file

Select the Service

☑ VIDEO IPTV
☑ VIDEO CDN
☑ VIDEO DTH
☐ SMART M2M
☐ M2M Jasper
☐ SMART GlobalSIM
☐ Internet of things
☐ SDPBR
☐ SDPMX
☐ Cloud Service Broker
☐ Consumer Electronics
☐ Insurance Telematics
☐ Open Web Device
☐ Partner Connect
☐ RCSD
☐ Instant Servers
☐ Security Hub
☐ Risk Management
☐ DSN
☐ Credit Card Broker
☐ Gnoc IP
☐ UDO

Select the Severity

☑ Critical
☑ Major
☑ Minor
☑ Slight
☑ Cosmetic

Select the Status

☑ Active
☑ Canceled
☑ Closed
☑ Delayed
☑ New
☑ Restored
☑ Solved

Select year

2016

Run with time ranges

NO

Select the Type

☑ Claim
☐ Order
☐ Query

Attention!

Are you really sure input controls are correct? Then, select the Generate Box. When the file had been generated, you can download it in the Download Button

Figure 2: Web interface to select the parameters of the data for the analysis.

Results

Due to a previous process mining analysis carried out for the years 2015 and 2016, we achieved that the number of ‘Major' incidents dropped and that restoration time became half from what it used to be. In 2017, the number of ‘Major’ incidents recorded in the trouble ticketing tool has still been as low as it was in 2016.

By the end of 2016 we started the search for non-evident operational drifts using ‘Severity’ changes as an indicator. The IoT managed connectivity service was one of the services with a higher percentage of change.

The process maps in Figure 3 show the changes of ‘Severity’ during the incident management for this service in 2017. The map includes all changes, no matter what their initial or final severity is, showing a complex flow.

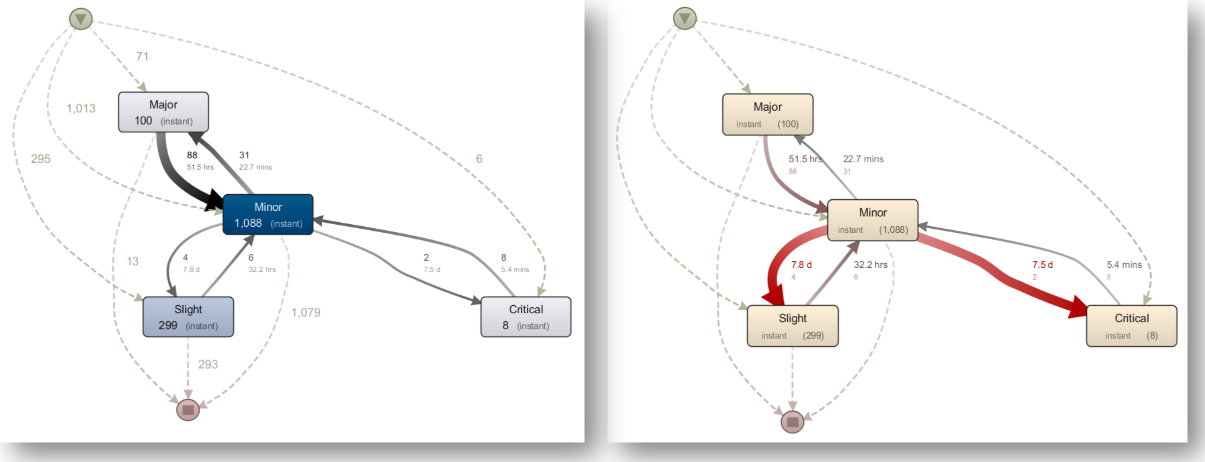

Figure 3: All changes of 'Severity' - Case frequency and Median duration - 2017.

Overall, service availability depends on 'Critical' and 'Major' Incidents. So, the first step was to use the Disco filtering to focus on the life cycle of these 'Severity' attributes and the changes that they went through. Figure 4 shows the changes of 'Severity' for incidents that started as 'Critical'.

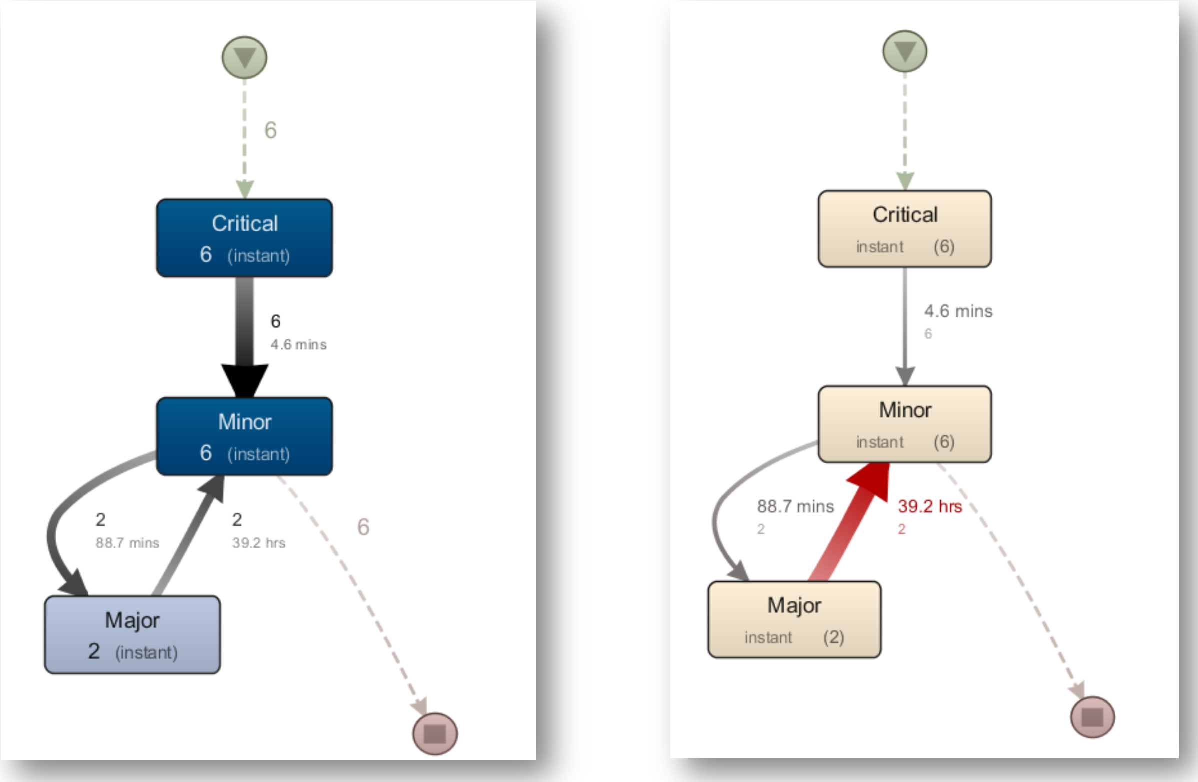

Figure 4: Change of 'Severity' in initially 'Critical' incidents - Case frequency and Median duration - 2017.

During 2017, there have been no final incidents of ‘Critical’ severity, although six incidents were initially issued with that level. They all were downgraded to ‘Major’ in less than five minutes. This is a common scenario. The issuing group detects a severe impairment and issues it so. By definition a ‘Critical’ incident has global impact. The operational group checks if the failure meets that condition. For instance, there may be a connectivity issue in country A, but the service keeps working in the rest of countries. In this situation, this is not a global service breakdown and the incident is reclassified as ‘Major’.

This practice is according to the protocol, so we diagnosed that ‘Severity’ management was correct for these six tickets.

‘Major’ incidents do not disrupt the global service but are taken very seriously by the operation teams. For instance, a ‘Major’ incident could be disrupting the business of a big customer, so the SLA times are very tight for this service.

In the analysis carried out for 2016, there was a notable decrease in ‘Major incidents’ compared to those recorded in 2015. If we add the information available up to the date of 2017, we note that this low level is maintained (see Table 1).

Period	2015	2016	2017
# Major Incidents	124	37	13

Table 1: Number of ‘Major’ incidents registered 2015, 2016 and 2017.

On first sight, this result indicates a remarkable improvement in the operation of the service. However, the high percentage of changes in ‘Severity’ led to a detailed analysis of what was happening (see Table 2).

	2016	2017
Opened as Major	177	71
Closed as Major	37	13
Change in Severity	144 (81%)	62 (87%)

Table 2: Number of incidents initiated as ‘Major’ vs. closed as ‘Major’ - 2016 and 2017.

The percentages of these ‘Severity’ changes were abnormally high (above 80%) compared to other services, which were in the range of 10%–20%. We did not perform hypothesis testing as the difference was so evident, but in order to automate detection a Chi-Square test is being applied in the future.

We performed a detailed analysis with Disco to detect at what point in the life cycle the ‘Severity’ change happens. For this, we combined both the ‘Updated Severity’ and the ‘Status’ field into the activity name (see Figure 5).

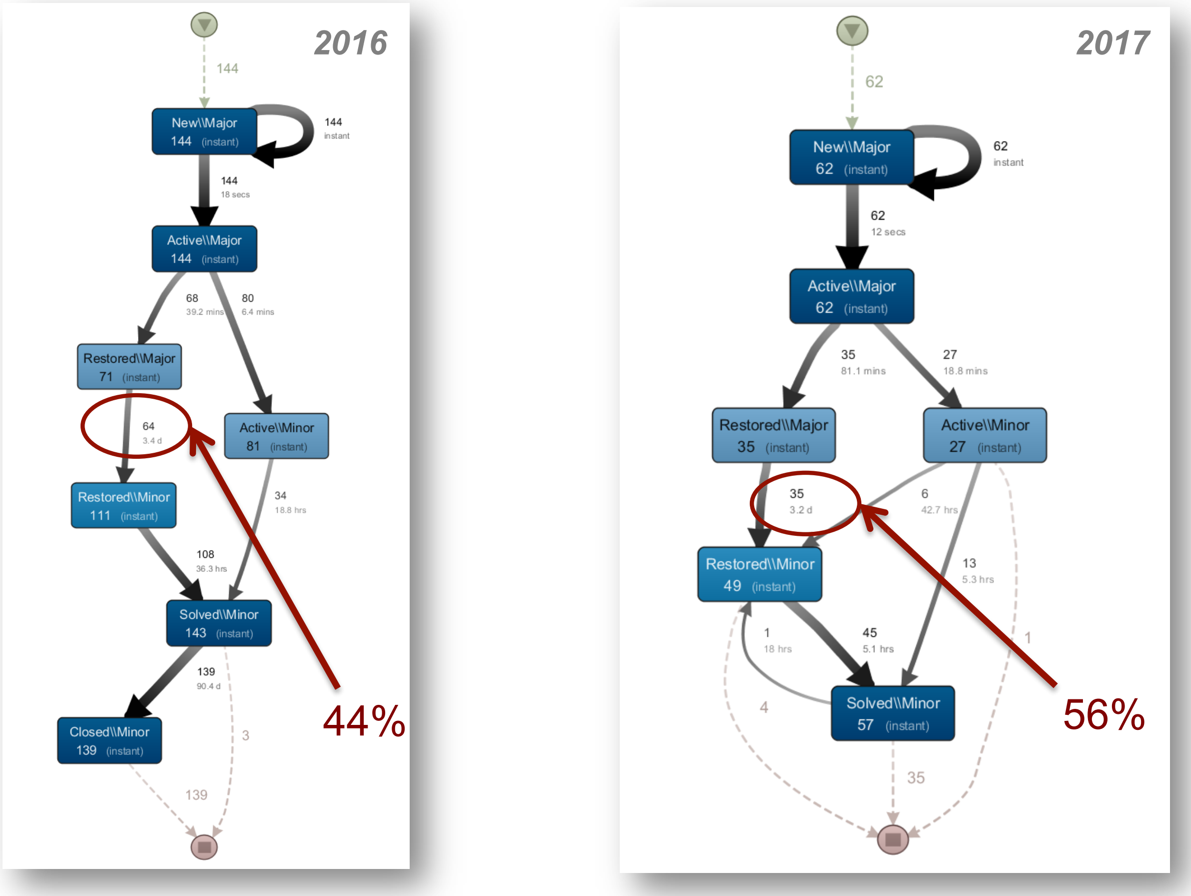

Figure 5: Status during the change of ‘Severity’ in ‘Major’ incidents - 2016 and 2017 (Case frequency and Median duration).

During 2016, there were 177 incidents initially issued as a ‘Major’ incident. 144 of them changed the ‘Severity’ attribute to ‘Minor’.

In a more detailed analysis, we found that 80 incidents changed the ‘Severity’ when they were in ‘Active’ status and only 20 minutes (median duration) after being opened. This is the typical downgrading of a misclassified ticket, an action that is normal and well-described in the protocol.

However, it is noteworthy that 64 changes in ‘Severity’ took place 3.8 days (median duration) after the incident restoration. This is an unexpected finding, because operational resources were spent fixing these incidents as ‘Major’. If the incidents were in fact not ‘Major’ then this would have been a waste of scarce human expertise. 44% of the total changes happened at this point (see left side in Figure 5).

In 2017, the behavior in ‘Major’ incident management has been similar to the one that we observed during the last year. 71 ‘Major’ incidents were registered and 35 changes in ‘Severity’ took place 4.6 days (median duration) after the incident restoration, which amounts to 56% over the total of changes (see right side in Figure 5).

In almost all cases, the operator that reduced the ‘Severity’ at this point belongs to the skilled global ‘Service Management Group’ (see Figure 6).

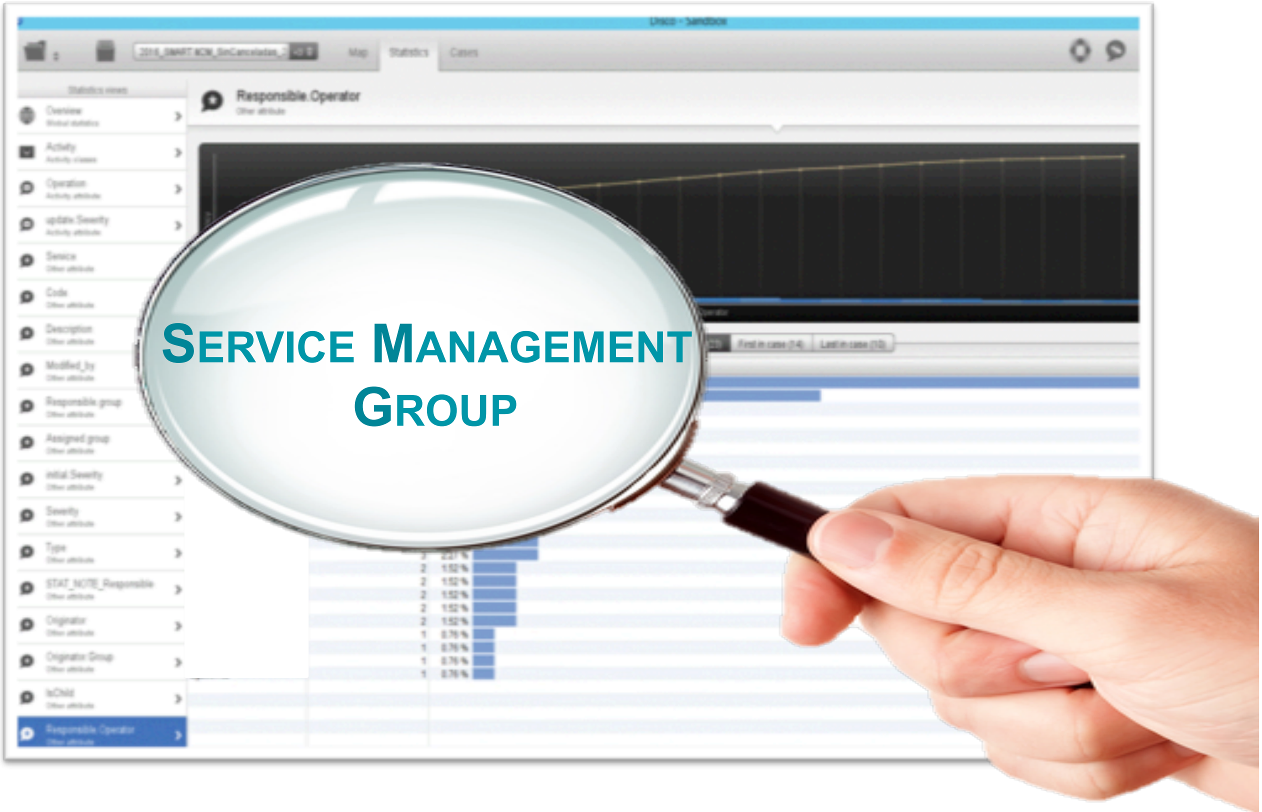

Figure 6: Operator responsible of the change of Severity on Major Incidents during 2017.

The Service Level Agreement for ‘Major’ incidents in this service is twelve hours. Looking for a clue to find the reasons of this behavior, we further analyzed the life cycle of the incidents that changed their status from ‘Major’ to ‘Minor’ once they were restored (See Figure 7 for the 2017 incidents where this happened).

The median duration to restore these incidents is 89.7 minutes, much lower than the twelve hours committed in the agreement with the client (see Figure 7). The question now was: “What is causing this unexpected behavior?”. The operations group has a high level of expertise, so a lack of information or training is not a reasonable explanation.

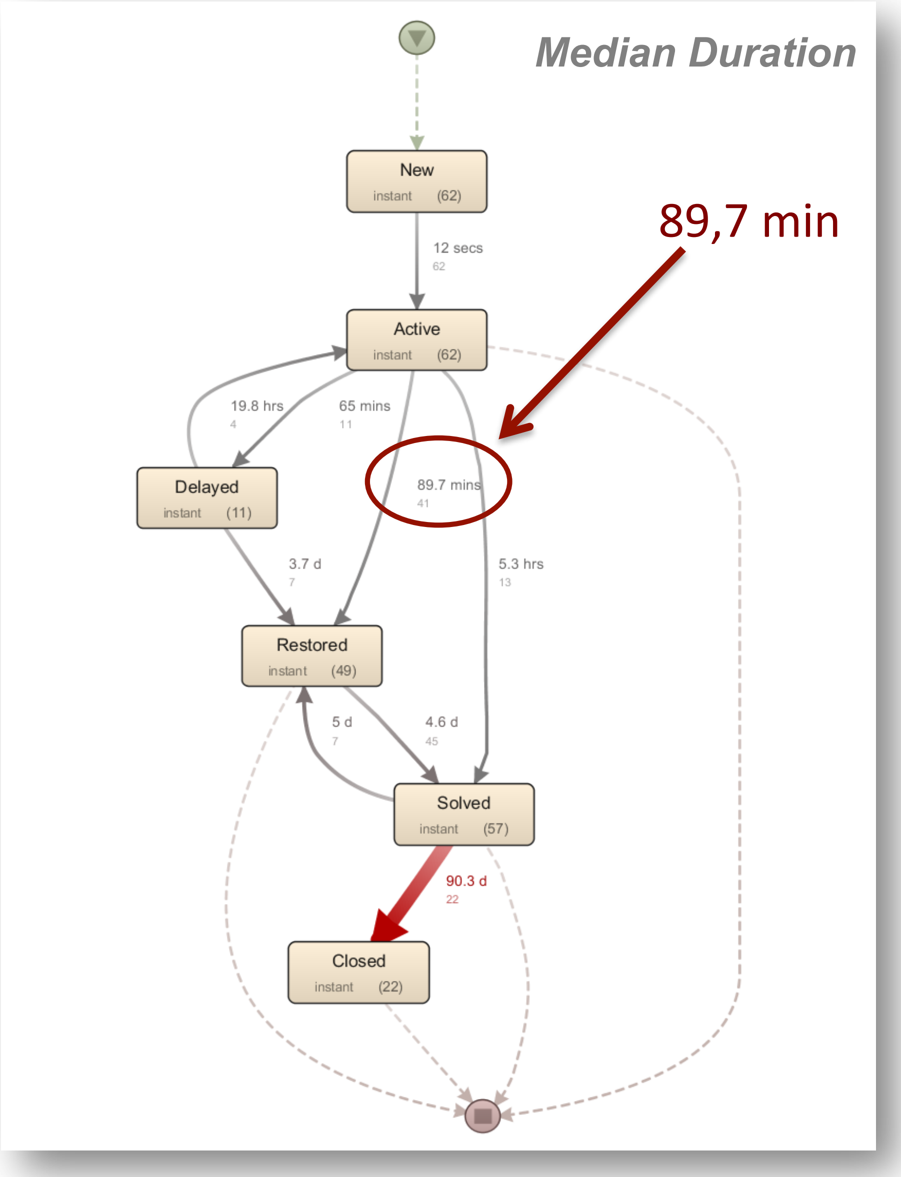

Figure 7: Life cycle of ‘Major’ incidents that changed ‘Severity’ in the ‘Restored’ status (2017).

Together with the operations team we found the source of these actions: There are ‘Major’ incidents that have no direct impact on customers but they must be restored with urgency to avoid a potentially bigger problem. For instance, the proactive surveillance detects a machine that is near to reach a dangerous level of CPU usage or a hard disk is becoming too full. These conditions do not affect customers at the moment, but once they have been detected they have to be fixed as soon as possible.

However, when these incidents were closed as 'Major' then they were recorded as such in the monthly reports — although they did not actually have any impact on the final customers. The management teams received the report and they raised concerns about the high number of incidents that disturbed their customers.

To avoid the reporting side effect for incidents, where this concern was not justified (because the incidents did not actually affect the customer), teams had started to reduce the 'Severity' of these tickets once they were restored. This is a good example of a hidden, informal cultural drift.

Impact

Based on our analysis, and the new understanding of why this unexpected behavior was happening, we identified the following two main problems:

1. Incidents were not registered correctly in the trouble ticketing tool. In this tool, it is actually possible to indicate whether a country is affected or not.

2. With the unofficial workaround, the information that was shared with the clients was reliable but the internal information was incomplete because many 'Major' incidents were not accounted as such.

As result of our process mining analysis:

- A hidden cultural drift in incident management has been fixed. This practice was hiding the real state of the service.

- Operational reports have been modified in order to reflect actual availability without the need to "manipulate" the recorded information.

- We have improved the relationship with our customers by providing more accurate and reliable information in our reports and communications.

'Severity' change analysis is now part of the monthly operational reports. If service managers detect an abnormal variation in the percentage of changes they may request a further analysis with Disco.

Authors

Javier García Algarra, Operations Data Analysis Leader (corresponding author)
Carmen Lasa Gómez, Data Analyst