Lessons Learnt Working with Performance Data in Call Centres

Tommaso Colombino, Benjamin Hanrahan, Stefania Castellani

Xerox Research Centre EuropeXerox Research Centre Europe
6 chemin de Maupertuis, 38240 Meylan, France
{firstname.lastname@xrce.xerox.com}

Abstract. This paper details the treatment of performance data in outsourced call centre operations, as encountered by a team of researchers throughout the course of a project. This project aimed at improving support for performance and motivation management in an outsourced customer care contact centre for a large telecommunications company. In particular, we focus on how the practices of capturing, aggregating, and presenting data reflect the operation's overall concern with "reporting upstream" and accountability. As well as, how the technological and organizational infrastructure of the call centre is shaped accordingly. We then discuss some emergent consequences of this organization of data management, which in particular take the form of some tensions between the emergent needs for data at certain levels of granularity and aggregation within the actual operations of the call centre, and its relative accuracy and availability.

Keywords: ethnography, ethnomethodology, call centre work, performance data management

1 Introduction

Customer contact centres, or call centres, provide an interesting but challenging context for technological innovation. They are knowledge work environments built around a fairly strict productivity model and division of labour. They have complex data and information management systems that, particularly in outsourced operations, are subject to a tension between the need for standardization, efficiency, and replication of infrastructures with the particular requirements of multiple client organizations. Consequently, a lot of attention has been devoted to modeling call centre activities and for making simulations, what-if analysis, forecasting trends, to schedule work, and to organize appropriate trainings for call centre employees [1, 4, 5, 6]. Some attention has been devoted to the challenges of innovating certain aspects of customer contact work and remote assistance [7, 9] and to the collection of insights that could contribute to a successful implementation of Customer Relationship Management system projects [10].

In our work we have focused less on the customer facing aspect in the work of individual agents (e.g. the aspect of problem solving over the phone) and more on the operational management of the call centre workforce as a whole, with particular focus

on the pervasive use of productivity metrics which are the direct product of the agents' call taking activities. In particular, this paper details the experiences of a group of researchers throughout the course of a project within an outsourced customer care contact centre for a large telecommunications company. The initial engagement with the call centre was through a series of ethnographic studies aimed at understanding performance management and compensation practices within the organization. In the course of the project we uncovered a number of emergent problems in the operations of the call centre. In particular, we focused on the lack of access to up-to-date information about performance and compensation for the agents of the call centre (the employees answering the phone and responding to customer queries) [3]. During our initial investigations we concluded that the primary causes for this lack of support were to be found in the lack of investment in innovation on the part of an organization built around a labour arbitrage business model, and a consequent lack of investment in the technological infrastructure of the call centre and the professionalization of its agents. To improve support for agents, we prototyped and deployed technologies designed to pull performance data from the call centres' existing technological infrastructure and provide near real-time feedback to agents and supervisors [2].

Through the process of integrating our technology prototypes with the infrastructure of the call centres, we discovered that the problem of data availability, accuracy, and relevancy (by which we mean aggregated and visualized to suit the purposes of a particular task at hand) was far more pervasive that we had initially assumed. This speaks to what Martin *et al.* describe as a more general problem of implementation of procedure and practice through technology [7], particularly how this problem becomes vexing in a Business Process Outsourcing organization. Mostly because the organization designs its technological and organizational infrastructure around the imperatives of reducing cost to minimum terms and demonstrating value to the client organization through standardized productivity measures.

In this paper we want to focus on what we learnt about the practices of capturing, aggregating, and presenting data, and how they reflect the operation's overall concern with "reporting upstream." We then discuss some of the emergent consequences of this approach to data management, with a focus on the tensions that arise from the prioritising of different data needs within the call centre, in terms of granularity, aggregation, accuracy, and availability.

2 Work Structure and Infrastructure of Observed Outsourced Contact Centres

The type of call centres that we focused on in our ethnographic studies are outsourced customer contact centres for large telecommunications companies (the client organisations) who provide standard customer account management services (billing inquiries, adding or removing services from the contract, etc.) and often also some level of technical support (troubleshooting software and hardware problems). Incoming calls are managed through a phone switch that receives calls from one or

more queues (which typically correspond to options selected by a customer through an Interactive Voice Response system) and distributes them to available agents.

Individual agents are logged into a phone switch that determines their "state" through a series of auxiliary codes (some of which the agents can select manually to put themselves in a particular state). An agent can therefore appear as "available" in the system and in a queue to receive incoming calls on a first-come first-serve basis, or "on a call," "on hold," "on a break," as having transferred a call to a different number, etc. Throughout any given shift in the call centre, all phone activity is logged, which means that there is a record of how much time any agent has spent in any given state and when.

This is one of the core mechanisms through which the call centre generates the data and the metrics, which are used to measure and report on its operational activities, and are generally referred to as Key Performance Indicators (KPIs). So, for example, the time the agent spends on the phone in conversation with a customer ("Talk Time") may be combined with the time the agent spends writing up the outcome of the phone call after it is finished ("After Call Work") and the time the agent keeps the customer on hold ("Hold Time") to create a metric that represents the overall time needed to handle that customer on that occasion ("Handle Time"). This metric can then be averaged over the time of a work shift, or week, or month of answered calls ("Average Handle Time").

The performance expectations of an outsourced call centre are typically defined by a Service Level Agreement (SLA), which is part of the service contract signed with the client organization. The call centre as a whole is expected to keep their aggregate average KPI values within the thresholds (or upper and lower threshold values) specified in the SLA – agents are therefore in turn expected to manage their phone calls so that their own average values fall within those thresholds as well. Typically, individual agents are grouped into teams of 10 to 15 workers to which supervisors are assigned from whom they receive periodic (weekly and monthly) feedback on their performance. Supervisors in turn report to Operations Managers, who are responsible for groups of 10 to 15 supervisors. This organizational hierarchy is designed to ensure control and supervision of a “floor” that may total 800 to 900 agents in some of the larger call centres.

Apart from the telephone switch data described above, the call centre also generates a large amount of "qualitative" data, which pertains more directly to the content of the interactions between agents and customers over the phone. The audio of all phone calls coming into the call centre is recorded. On top of that, a video screen capture of the agent's desktop is recorded for a subset of calls along with the audio. A sample of the total calls recorded for each agent is listened to - either by a supervisor, a quality assurance manager, or a member of the client organization - coded, and scored to produce a Quality Score. A small subset of the customers who call on any given day are also called back and asked to respond to a Customer Satisfaction Survey (CSAT). Like the phone switch data, Quality Scores and CSAT scores are initially associated to individual agents and then aggregated to produce overall performance metrics for the call centre as a whole.

To give an idea of the overall amount of data that is generated, we can say that in the type of call centres we observed, an agent may take up to 60 - 70 calls in an 8 or 10 hour shift and that the number of agents for a single call centre (which may only

be one of many operating on behalf of any given customer) is in the hundreds. And it is important to understand that while the core mission of a call centre is to provide assistance to customers, the logged data, that is the by-product of the interactions between customers and agents, is what is at the heart of the interaction between the client organization and the service provider, and having and presenting the right data is therefore a key operational concern.

3 How is Data Used in the Call Centre?

The actual technical infrastructure through which data is managed can vary slightly, but typically an outsourced services provider that manages customer contact centres will have a centralized Enterprise Data Warehouse. This warehouse collects all of the phone switch and qualitative data from various centres. The source of this data is a point that varies as well.. FFor some functions within the call centre the data is pulled directly from the phone switch, and for others it is pulled from the data infrastructure of the client organization.

In the case of data coming from phone switches, it is important to note that the time logs of agent states or activities assigned to individual auxiliary codes are the raw data from which the Key Performance Indicators (KPIs) need to be calculated (as describe above in the case of Average Handle Time (AHT), which is the sum of a number of individual auxiliary codes divided by the number of calls over an interval of time). This means two things: (a) data as it is produced is not necessarily in itself "readable" or inherently meaningful; (b) given the sheer amount of data involved, producing readable data at any particular level of granularity or aggregation that might be relevant for a specific person or purpose at a specific point in time is not trivial in terms of time, computational cost, and infrastructure requirements.

In the point of view of the outsourced call centre provider, these and other factors lead to a fluid, non-homogeneous environment that is difficult to support with a general, technical solution for all customers. This results in the creation of many ad-hoc processes, e.g. reporting the current KPIs that agents need to focus on (these can change from month to month), checking the adherence of the agents' tracked status to their individual schedule (there are many exceptions that change from site to site), and upstream reporting. For all the flexibility that these ad-hoc processes provide, they lack scalability, reusability, and automation, often resulting in a fair amount of work to maintain and adapt to new needs.

The amount of work involved in these ad-hoc processes is tolerated (and often considered necessary) because periodic performance data is the key output for the service provider. This is how they communicate to the client organization how successfully they are in fulfilling the Service Level Agreement and often the basis for how the call centre is compensated for their work.

Based on our observations of call centre work and management, we can describe three fundamental uses of the data, each of which has its own criteria to meaningfully aggregate and represent or visualize the data:

- Real-time or near real-time monitoring of agent activity and overall operational status of inbound call activity;

- Agent performance management and feedback (including performance reviews, coaching, and calculation of compensation);
- Upstream reporting at various levels (agent team to call centre leadership, call centre to central operational management, service provider to client organization).

Real-time Monitoring of Agent Activity and Metrics

The real-time or near real-time monitoring of phone switch and other types of data is done mostly to identify emergent and potentially problematic outlying situations. An example of this would be an agent who has been on a call for a very long period of time or not at their workstation when they should be.

In order to monitor for any of these situations, managers within the call centre typically have access to live or near real-time visualizations of switch data, in the form of tables that display the switch data for each agent and each auxiliary code that is being logged. With these tables managers can spot any obviously outlying situations. However, it should be noted that due to the large variation in what is a 'problematic' situation (in terms of both the values for a similar metric or even the metric itself), the support for locating these situations is minimal. Managers are required to do a visual scan of the table periodically. This ad-hoc approach to detecting 'problematic' situation is an example of a hidden inefficiency and creates new organizational issues. Most notably, there is a need to determine how much time should be devoted to the real-time monitoring of switch data and by whom. We observed several cases of agents who were on a particularly long and difficult phone call asking for assistance without waiting for someone to notice the abnormally long talk time on the switch monitor. Another aspect that should be noted is that this live switch data does not provide a longitudinal view or aggregate value in any way, and therefore is not useful in analysing performance trends.

Another example of near real-time monitoring can be found in the case of a metric that is generally referred to as "schedule adherence." Schedule adherence represents the adherence of the agents' tracked status, as detected through the phone switch, to their individual schedule. To give an example of how this works: an agent is scheduled to take his or her 10 minutes break at 11:00 am, but they are on a call with a customer until 11:02. At 11:02 they finish the call and put themselves in the auxiliary status for "break." They take their 10 minutes break and come back at 11:12. At this point they are 4 minutes out of adherence (2 minutes at either end of the 10 minutes break). This can be calculated automatically by matching the logs from the phone switch with the "expected" values as entered in a scheduling system. What cannot be detected automatically is whether any time "out of adherence" can be considered legitimate or not (the terms of which are usually defined in the SLA). In the case of the example above, being on a call with a customer may be considered a legitimate reason to deviate from schedule, and the 4 out of adherence minutes can be treated as an exception that will not be counted against the overall adherence metric.

An interesting feature of this is that in many of the call centres we observed there was no established or firmly enforced procedure for communicating exceptions, it was mostly or entirely done ad-hoc. They can be communicated by e-mail on a piecemeal basis, or collected on a paper roster by a supervisor who then e-mails a list of exceptions to workforce management, and keeping track of exceptions effectively

depends on how prompt agents are in communicating them to their supervisors, assuming they remember to do so in the first place. One net result of this is that workforce analysts actually spend a considerable amount of time chasing down unreported exceptions. This requires a manual examination of adherence reports for individual agents, which display a time-stamped sequence log of agent auxiliary codes and whether each entry in the log is consistent with the schedule or not. Unreported inconsistencies between the log and the schedule can then be chased down with agents and supervisors to determine whether they fall under the terms of the SLA or not. This is a time consuming process which is difficult to automate because not all legitimate exceptions correspond to a specific auxiliary code - it is common practice, for example, to have generic auxiliary codes for a number of "unavailable" or "away from work-station" conditions which do not need to be differentiated for reporting purposes.

Agent Performance Management and Feedback

Agent performance management and feedback has two broad goals: to provide agents with the right level of "situational awareness", i.e. an understanding of their ongoing performance within the broader context of the operations of the call centre as a whole; and to provide the right motivation and performance related incentives to employees. Coaching sessions between agents and their direct supervisors typically take place every week, when supervisors review the ongoing performance of an agent (both with respect to the qualitative and quantitative metrics they are assessed on) and may provide specific goals and objectives for the agent to attain, as well as practical advice on how to get there.

Agents themselves generally have limited independent access to up-to-date and detailed data about their own performance, and we have addressed this issue and its consequences in some detail in a previous publication [3]. They do typically have access to a scorecard like the one shown in Figure 1.

Figure 1: Agent Performance Score Card with KPI values displayed on a rolling 30 day average

The scorecards give the agents an overview of their KPIs at a level of aggregation, which may or may not be relevant from the agents' point of view (for example, for the purpose of understanding how their compensation will be calculated or how far they are from achieving a performance based compensation target).

There are multiple reasons for the call centre to limit agent access to data, even if that data may very well be relevant to their everyday requirements. In the first instance there is a cost in terms of infrastructure and development that may be difficult to justify in a business model that is basically grounded in labour arbitrage and optimization of overhead costs. Part of that cost may actually be tied to licensing fees required to access and visualize switch data on individual workstations. On top of that, many call centres, rightly or wrongly, do not want to provide agents with easy access to data that may either distract them from their core activity of managing customers over the phone or indirectly encourage them to "work to the numbers" – i.e. leverage their understanding of the performance assessment and compensation mechanisms in order to hit the highest compensation rates without necessarily providing the best value to their organization [3] – if not outright game the system. In other instances the availability of performance data is controlled by the client organisation and the service provider can only pull the data at a predefined interval, e.g. once a day.

Whatever the reasons for limiting access to performance data within the call centre at the agent level (by which we mean both that the data is not directly accessible to the agents, and that where it is accessible it is not necessarily aggregated and visualized at the relevant level), the fact is that there are emergent needs for that data, and as we shall discuss in the following section, those needs are often addressed using data and systems that were put in place for different purposes.

Upstream Reporting

What should emerge quite clearly from the previous discussion of the organizational set-up of an outsourced call centre is that the measured performance of the call centre as a whole is quite literally the aggregate performance of all its agents. And given the number of levels in the organizational hierarchy (from agent, to team, to operations management, to call centre management, to multi-site client account management), it was no surprise to learn that there is quite a lot of reporting going on, and most of it upstream. Putting aside the agents themselves, who are at least spared the burden of having to create reports of their own performance by their supervisors (although as we have seen this also means they do not typically have access to data about their own performance), everybody else has to produce regular reports about their own area of responsibility.

This is the mechanism of performance management and accountability in the call centre – data is progressively aggregated and pushed upstream so that each level of the management chain can demonstrate that they are meeting their own performance targets to the managers above them, and hopefully performance issues that may have an impact on the call centre as a whole are caught and managed somewhere along the way.

It is a mechanism that also indirectly addresses a potentially costly problem for the call centre. In this paper we have discussed at length how telephone switch data is handled, because it is critical and because it is highly present in the call centre's

everyday operations. But switch data is by no means the only kind of data – we have also talked about quality assessment data, customer survey data, scheduling and attendance data, payroll data, etc. All these different types of data are captured, stored and visualized through a number of different information systems, some perhaps proprietary to the client organization, some legacy, some off the shelf, some developed in-house by the outsourced service provider. Integrating these different sources of data at the infrastructure level is a cost that most outsourced service providers cannot afford. It is expensive as a development project in itself, and it would require costly and frequent adaptation and modification to provide everybody access to the requisite, integrated data.

Distributing the responsibility and the task of aggregating the data for particular purposes across the various levels of management means that the organization can get away with not having an integrated solution for handling all its data requirements. However, it also means that where data requirements emerge that had not been anticipated, or are not considered strategically important enough to devote an organizational process or information system to the task, people have to come up with their own solutions and workarounds.

In the following section, we will provide some examples of emergent data requirements, the sometimes less than ideal compromises used to meet those requirements, and ultimately the impact this problem had on our attempts to engage the business in research and innovation.

4 Practical Compromises in Data Management and Availability

In the previous section we outlined several of the ways in which data was designed to be used in the call centre. Within these designed uses, there were still aspects that were ad-hoc due to the fluidity of business requirements and non-homogeneity of the different call centres operated by the service provider. In this section we outline some of the emergent appropriations of these systems to solve new emergent problems within the call centre.

Detecting Agents Who Cheat

Live switch data can also help call centre management spot agents who are attempting to "game" the system to their advantage. An example of this was observed in call centres that pay their agents at hourly rates, regardless of the actual time spent on the phone with customers. The way agents would "game" the system was to continuously switch themselves in and out of available status to receive a phone call. Every time they did this, the phone switch would place them at the back of the queue to receive incoming phone calls. If used judiciously, this trick could effectively allow agents to go through an entire shift without taking any call whatsoever. The interesting feature of this example is that it would be relatively easy to detect it "indirectly" by looking at different representations of the data. An agent that takes no calls over a shift will have an average handle time that will not even compute for that shift, as it would have to divide by zero calls. However, an agent that takes one or two calls over a shift may have a relatively normal looking average handle time value for

that shift, in which case the revealing metric would be the number of calls per hour or total calls taken over the shift.

This, while seemingly straightforward, reveals itself to be problematic in practice for a specific organizational reason, which has a direct consequence on the way the relevant data is routinely managed and made available: a call centre that pays agents through a flat rate as opposed to some productivity based mechanism is less likely to focus on metrics that directly relate to any individual agent's productivity, or if they do, they are unlikely to look at them at the right level of granularity. In other words, they will not routinely generate reports that show average values for talk time or calls per hour over a single shift (more likely they will for a period of a week, or on a rolling 30 day average), and may not have easy access (not to mention the time and willingness) to the tools to generate a daily report for the relevant metrics.

Since it is not easy to detect this behaviour, we observed call centre workforce management staff appropriating the switch monitor to detect agents "gaming" the system. This system is not designed for this use and does not provide a "record" of agents with revealing performance metrics, it does however provide the opportunity to observe agents switching themselves in and out of the available queue in real-time. As this emergent need had not been designed for, the call centre workers developed an ad-hoc solution. This introduces the identical practical problem as before, since this ad-hoc solution requires constant monitoring of the system by someone - the trade off between time and effort versus gain becomes very difficult to measure.

Providing Agents With Useful Performance Feedback

As already mentioned, supervisors in a call centre typically manage a team of 10 to 15 agents. While they may have different roles or responsibilities (i.e. specialize in dealing with specific queues or categories of problems), the key reason for dividing the workforce up into teams is to enforce an appropriate level of performance management. Supervisors therefore have two key responsibilities: coach, monitor, and report on the performance of their agents; and provide assistance in handling calls and customer issues when necessary.

As we discussed earlier in the paper, the level of access that the agents get to their own performance data varies greatly from place to place. In some call centres, Information Management Systems available to agents may provide access to live switch data, which allows them to see for how long they have been in any particular “state” (e.g. on a call, waiting for a call, unavailable, etc.). Because performance targets for agents (and for the call centre as a whole) are expressed as average values over some period of time (which typically corresponds to a pay period), live switch data does not give agents intelligible feedback about their ongoing performance with respect to their goals. The problem is that the reports or scorecards (see Figure 1 for an example) with aggregated data the agents are typically provided with do not provide a very “dynamic” view of their own performance – they are updated irregularly, and always provide the same “view” of the data (for example, KPI values displayed on a thirty day rolling average).

Average KPI target values or thresholds for most agents do not change much from day to day, as they are written into the Service Level Agreement between the client organization and the outsourced service provider. However, there are two exceptions

to this that create emergent data needs that agents and supervisors have to manage themselves.

The first exception arises from the optimization of agents' productivity, some agents work on more than one phone queue, each one of which corresponds to a different category or level of assistance (for example, providing support for both mobile phone and desktop computer products). These queues may have different target KPI values, and therefore the agents may have individual KPI targets that "shift" dynamically depending on how many calls they take of each type. This has the effect to render whatever tools are available to agents useless, as they were not designed with this in mind and can only show one set of target values (causing wildly inaccurate data to be displayed). This creates the need for supervisors to improvise solutions in order to communicate performance metrics to agents.

The second exception emerges from the fact that agents, irrespective of what the general target values for KPIs are within the call centre, are often given interim goals by their supervisors on a weekly basis. These interim goals are typically agreed to in coaching sessions with supervisors and designed to provide realistic and attainable goals in the short term, in particular for agents who are underperforming.

In practice this means that while agents are usually aware of what their individual targets are, it is harder for them to know where they are, in their own performance, with respect to those targets at any particular point in time. From the agents' point of view, this is not an issue that directly concerns their ability to provide assistance to customers over the phone. But because the organization is concerned with keeping its metrics within the parameters defined by the SLA, the agents do experience a certain level of direct (through their supervisors) and indirect (through compensation mechanisms that are at least in part performance-based) pressure to keep their KPIs within the expected values.

It falls upon the supervisors therefore to ensure that the agents have the right level of awareness of their ongoing performance. As mentioned earlier, there currently is no information system dedicated to providing agents with relevant performance data in a timely fashion [3] in call centres. Moreover, the subset of KPIs that might be relevant to a particular agent or team of agents is likely to be spread across a variety of different information systems that capture the data from a variety of sources.

For example, in one call centre we observed that a supervisor needed to provide his agents with performance feedback for the following KPIs: value points (an arbitrary metric that assigns points that are used to calculate bonuses and provide a normalized ranking of agent performance); calls per hour (CPH); average handle time (AHT; customer satisfaction survey scores; and percentage of calls transferred. Of these metrics, three (value points, CPH, and AHT) could be pulled from a single information system the supervisor had ready access to, one (customer satisfaction survey scores) was available on a separate, client proprietary system, and the final one (calls transferred) was on a system the supervisor did not have access to at all, meaning he relied on his operations manager to forward the relevant data to him on a daily basis.

To make the data more relevant and accessible to the agents, the supervisor decided to create a series of custom pivot tables in Excel with custom formulas that cover each metric into a percent to goal value, colour coded red or green to provide

immediate visual feedback on whether the values are within the thresholds or not. Figure 2 shows an example of the Custom Pivot Table.

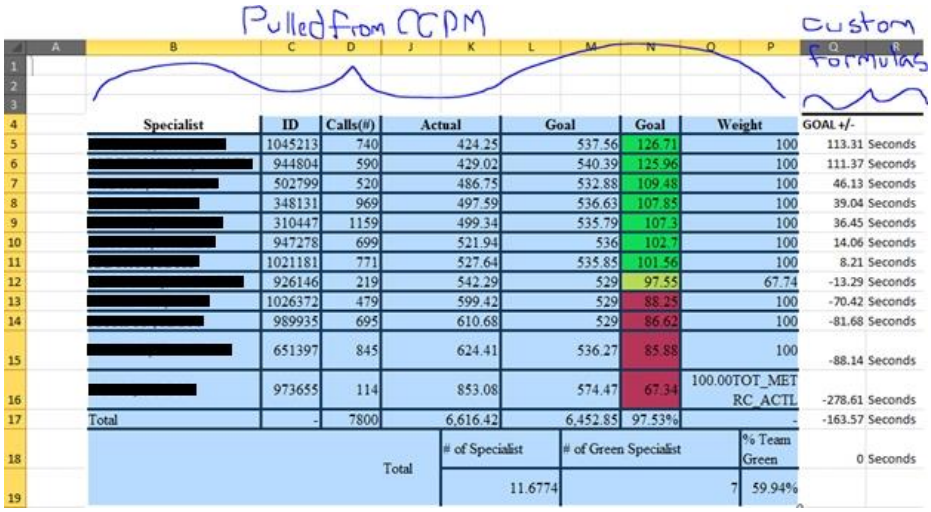

Specialist	ID	Calls(#)	Actual	Goal	Goal	Weight	GOAL +/-
[redacted]	1045213	740	424.25	537.56	126.71	100	113.31 Seconds
[redacted]	944804	590	429.02	540.39	125.96	100	111.37 Seconds
[redacted]	502799	520	486.75	532.88	109.48	100	46.13 Seconds
[redacted]	348131	969	497.59	536.63	107.85	100	39.04 Seconds
[redacted]	310447	1159	499.34	535.79	107.3	100	36.45 Seconds
[redacted]	947278	699	521.94	536	102.7	100	14.06 Seconds
[redacted]	1021181	771	527.64	535.85	101.56	100	8.21 Seconds
[redacted]	926146	219	542.29	529	97.55	67.74	-13.29 Seconds
[redacted]	1026372	479	599.42	529	88.25	100	-70.42 Seconds
[redacted]	989935	695	610.68	529	86.62	100	-81.68 Seconds
[redacted]	651397	845	624.41	536.27	85.88	100	-88.14 Seconds
[redacted]	973655	114	853.08	574.47	67.34	100.00TOT_MET RC_ACTL	-278.61 Seconds
Total	-	7800	6,616.42	6,452.85	97.53%	-	-163.57 Seconds

Total	# of Specialist	# of Green Specialist	% Team Green	
	11.6774	7	59.94%	0 Seconds

Figure 2: Custom Pivot Table showing agent Average Handle Time values and % to goal

The supervisor was then able to share the tables with his team, either by forwarding them through an internal chat client (if he felt it was urgent to bring the metrics to the attention of the agents), or by posting them on a shared team home page for the agents to access at their discretion.

While what was described above may seem like a relatively routine operation (and well within the expected responsibilities of a supervisor), it nevertheless is a less than ideal compromise from the supervisor's point of view. It leaves each agent at the mercy of their supervisor for information on their performance. More critically, the entirely manual operation of pulling data from various systems, pasting it into a spreadsheet and forwarding the information to the team, only creates a static snapshot of the data at a particular point in time, when the goal in fact is to provide agents with feedback about their *ongoing* performance. In other words, the entire operation needs to be repeated every time a supervisor wants to provide agents with an update on their ongoing performance.

Introducing Technological Innovation

These different observations about the data needs of the call center, particularly the deficit in communicating performance goals and progress between agents and their supervisors, led us to propose an intervention. The two goals for our system were to give better performance awareness mechanisms to agents and to give a more responsive channel of communication between agents and their supervisors.

Our initial effort was in providing the agents with a responsive, real-time visualization of their current performance with respect to their KPIs. This is detailed in a previous publication [3]. Our follow-up effort was to provide supervisors and

their agents with a flexible infrastructure through which they could launch games and competitions that could be more responsive to both the needs of call centre and the more fluid needs of individuals [2].

One of the key objectives of our project was to provide much more responsive data to agents and supervisors. To realize this goal and deploy a tool to accomplish our goals required a significant amount of interaction with both the IT staff of various call centres as well as the developers for the service provider. This was necessary to understand how our tool would fit into the overall ecosystem and what capabilities needed to be developed by all interested parties. Trying to fit our tool within the existing infrastructure and obtaining the data with the required accuracy and timeliness that we required for our real time applications, turned out to be next to impossible. In some cases the infrastructure could not handle the load of providing real time reports of the data. In other cases the data was simply not made available by the client organizations in anything approaching real time. In the cases where our tool eventually was deployed the amount of effort required of the IT staff to provide near real time data was in the order of six to eight months.

Through these interactions with this arm of the organization and the detailed interrogation of the infrastructure through trying to deploy within it, we gained a much deeper insight into the reasons for the numerous improvised solutions that we encountered. We gained insight into how the data takes time to validate due to the imperfect and sometimes contradictory nature of the instruments used to capture the data. How the data's accuracy was negotiated and validated between the service provider and their client; which in turn impacted the rhythm that it was made available; and how changes to the different systems were prioritized. We often found that the IT staff and development staff were running at full capacity to keep up with the small set of emergent needs that arise for just communicating around the SLA with the client.

Through this process we learnt more about the call centre operation and policies than we did about the efficacy of our tool. The lessons that we learnt often represented limitations and barriers. Interestingly, and informatively, these are the same limitations that the call centre agents, supervisors, and managers also operate under. Ultimately, where our experiments and deployments were somewhat successfully completed, it was mostly because we found ways to improvise solutions to the barriers and limitations that we encountered, rather than by directly addressing the weakness in the technology or infrastructure.

5 Conclusion

In our initial investigations of call centres we observed many ad-hoc solutions to emergent performance data management problems in call centres. It was when we ourselves had to interrogate the infrastructure in which workers encountered, developed, and used these tools that we came to understand more about why they existed, and how addressing them was unlikely to be a simple matter of introducing technological innovations targeted at the problems revealed by the ethnographic studies. The way in which the emergent needs of workers and their supervisors are

handled, sometimes in spite of the infrastructure, illuminates their problems and priorities. If an organisation believes that developing an ad-hoc solution, be it a tool or process, to a problem is merited and takes the time to build it, this is significant. The overall infrastructure that is provided gives insight into how management perceives the work of agents and supervisors and how they weigh these needs with their needs for upstream communication.

Thus, the infrastructure is a clear representation of the priorities of the organization as a whole. In our experience, it seems easy for management to forget that the lower priority items of workers must be given a degree of priority as time progresses or this functionality will suffer from starvation of resources. For these call centres the infrastructure that was developed shows a clear imbalance of functionality and support for upstream communication. This is understandable when situated within our other knowledge of call centre workings, as service providers have an important obligation to clearly provide upstream communication to the client organization.

Through this lens we can now provide answers to questions that vexed us at the beginning of the study. Data, that according to common sense should be captured, stored, and be made available in real time is not necessarily available in real life. This is because this capability is simply not important for the construction of reports for upstream reporting, nor is it critical to solving customer needs. It is however, critical information to provide to workers to understand their performance within the constraints given to them. Data is not often in the format one needs for your application or emergent need since they are logged according to the very specific needs of different applications.

Robust, reusable solutions to capturing data simply do not exist in this organization due to a business model based on labour arbitrage and low margins. Data is not always correct and sometimes even clearly wrong, but this is not visible or deemed important into an application or process that makes use of them. Applications that may be useful to some actors, may not be authorized by managers because that would provide those actors with access to data that the managers do not want those actors to have access to. Sometimes this is due to an aforementioned problem of data accuracy, if a new application is now displaying and using data at an earlier time in the process this creates problems by having to validate the data earlier, sometimes even at a lower granularity.

Throughout the course of our technology innovation project with the outsourced service provider, we acquired a rich understanding of the situated practices of the operational part of the organization through observation. Where the practical constraints of the research project did not allow us to observe the whole organization (and indeed assuming that the priorities, implicit or explicit, that apply to infrastructure and technology development and deployment would be directly observable in action at any other level of the organization), a hands on interrogation of their infrastructure through prototype development and integration provided an additional level of insight into the relationship between the business model, the organizational hierarchy, and the situated problems experienced by the call centre agents and managers we observed. This insight may turn out to be more valuable than the technology products that were ultimately developed and deployed, which essentially are more technologically sophisticated workarounds to the ones we originally described. That is, this value will be realised only if the organization itself

gains awareness of the real impact of the strategic choices in technological infrastructure, research, and development. Defined and constrained as they may be by the business model and the routine work of its operations, they are also (perhaps unwittingly) constrained by the cost of those choices in terms of time lost, relative accuracy of its productivity measures, and ultimately quality of work experience for its employees.

References

1. Bourne, J., Murray, E., Iannone, J., Keren, S., McLean, N., Bourke, M., Fama, J., Watson, J., Nies, J. G., and Lverly, T. (2008) "SYSTEMS AND METHODS FOR WORKFORCE OPTIMIZATION AND INTEGRATION", US published patent application 20080181389, July 2008.
2. Castellani, S., Hanrahan, B., Colombino, T., and Grasso A. (2013). "Game mechanisms for production environments". In CHI 2013 Workshop on Designing Gamification: Creating Gameful and Playful Experiences, Paris, France, April 27, 2013.
3. Colombino, T., Castellani, S., Grasso, A., and Willamowski J. (2012). "Agentville: supporting situational awareness and motivation in call centres". In Proc. Of the 10th International Conference on the Design of Cooperative Systems, Marseille, France, May 29- June 1, 2012.
4. Cooper, K. (2008a) "Performance motivation systems and methods for contact centers", US granted patent 7,412,402, August 2008.
5. Ibrahim R., L'Ecuyer P., Regnard, N., and Shen H. (2012). "On the Modeling and Forecasting of Call Center Arrivals". In Proc. of the 2012 Winter Simulation Conference, Laroque, C., Himmelspach, J., Pasupathy, R., Rose, O., and Uhrmacher A. M., eds.
6. Mehrotra, V. and Fama, J. (2003). "Call Center Simulation Modeling: Methods, Challenges, and Opportunities". In Proc. of the 2003 Winter Simulation Conference, Chick, S, Sanchez, P. J., Ferrin, D., and Morrice, D. J., eds., pp. 135—143.
7. Martin, D., O'Neill, J., Randall, D., and Rouncefield, M. (2007) "How Can I Help You? Call Centres, Classification Work and Coordination". Computer Supported Cooperative Work, Springer. 16:231-264.
8. Martin, D. and Rouncefield, M. (2003). "Making the Organisation Come Alive: Talking Through and about the Technology in Remote Banking". Human–Computer Interaction, vol. 18, nos. 1 and 2, pp. 111–148, Lawrence Erlbaum Associates.
9. O'Neill, J., Castellani, S., Roulland, R., Juliano, C., Dai, L., and Hairon, N. (2011). "From Ethnographic Study to Mixed Reality: A Remote Collaborative Troubleshooting System." In Proc. of the 2011 ACM Conference on Computer Supported Cooperative Work (CSCW), March 19-23, 2011, Hangzhou, China.
10. Saeed, K. A., Grover, V., Kettinger, W. J., and Guha, S. (2011). "Organisational Interventions and the Successful Implementation of Customer Relationship Management (CRM) System Projects". The Database for Advances in

Information Systems, Volume 42, Number 2, May 2011.