

## Service Infrastructures: A call for ethnography of heterogeneity

ROGÉRIO A. DE PAULA

VICTOR CAVALCANTE

CLAUDIO PINHANEZ

*IBM Research - Brazil*

*This paper investigates the notion of heterogeneity, inspired by Latour's work on Actor Network Theory, as a lens for understanding daily work practices in a large service delivery organization. To this end, we present and discuss the findings from an ongoing research where we unpacked how system-administrators manage and negotiate incident resolution requests as part of service delivery practices. In particular, we looked into how performance metrics, such as, service level agreements (SLAs), mediated those practices. This paper contributes to the studies of infrastructure and explores the critical synergy between quantitative and qualitative methods in support of large-scale work practice research.*

### INTRODUCTION

Managing large information technologies (IT) infrastructures—vast configurations of heterogeneous actors and artifacts—constitutes the everyday service practices of a large-scale IT services delivery company. They form assemblages (and networks) of servers, network cables, routers, and protocols, work processes, operations, performance metrics, time zones, business interests and values, clients, system-administrators (sysadmins), salespeople, management, and practices, which are constantly being activated (and deactivated), brought together, and brought into effect as result of ongoing efforts to orchestrate and maintain the overall socio-technical system operational. That is, while customers infrastructures are taken as a stable and transparent substrate against which the actual business operations unfold (or they would seriously like to take them as such), for IT service organizations such infrastructures are manifest aspects of their everyday work routines. They are the locus and focus of IT service organizations' business concerns. This “infrastructural inversion” (Bowker 1994)—foregrounding the actual backstage elements of service delivery practices—helps to unveil and account for the ways in which those actors and artifacts come into play shaping (and being shaped by) the daily work of service delivery people.

This study thus looks into service delivery through the lens of infrastructure maintenance practices. That is, we are interested in the standard forms of representing, assigning, and reporting incident resolution that shape the logics and practices of service delivery in an organization. Following the Star et al.'s longstanding line of ethnographic research of infrastructure (Bowker et al. 2010; Star 1999; Star 2002; Star and Ruhleder 1994; Star and Ruhleder 1996), we set out to examine some of the more mundane aspects of service delivery. To this end, we set out to study the processes whereby not only are contractual agreements (i.e., service level agreements—SLAs) artifacts for measuring and ensuring high quality levels of service delivery, but also for negotiating a team's performance. That is, we examined the practices of incident dispatching (read, problem handling) and the ways in which they

## RENEWING SERVICES

are often modulated by a particular customer's SLA. We gathered narratives around the tensions between the local and the global, observed the ways in which existing (homebrewed or global, standardized or otherwise) tools affect delivery practices, and examined how incident 'tickets' become currency for performance assessment and negotiation.

In this paper, we thus attempt to explore and expand the Star et al.'s notion of infrastructures by introducing Latour et al.'s notion of heterogeneity (Latour 2005; Law 1992). In this context, infrastructures are not different from society in that they can be thought of as assemblages or networks of heterogeneous 'things' (animated or otherwise). In brief, Latour's Actor Network Theory (ANT) suggests that society, organizations, agents, and machines are all results of patterned networks of diverse/heterogeneous materials (Law 1992), be they people, machines, policies, organization hierarchies, and so on. As these materials interact with one another, they leave behind traces that we usually call "the social." While Star et al. recognize heterogeneity in their ethnographic accounts and theorizing of infrastructure and consider its relational property—i.e., it emerges out of its relations to organized human practices (Star 2002), they still separate the underlying material substrates from the actual practices. According to ANT, on the other hand, humans, boxes, networks, software applications, documentations are all actors in those heterogeneous networks (or assemblages). We believe then that this 'messier' account of service delivery 'infrastructuring' can help shed new light on how different network configurations of materials come to affect service delivery practices, performances, and results.

## STUDIES OF SERVICE INFRASTRUCTURES

Although IT outsourcing service is a major, business critical, and growing industry worldwide, ethnographic studies of everyday practices inside service delivery organization remain still few. There are remarkable exceptions, notwithstanding. Back in 1997, Sandusky (Sandusky 1997) investigated the work of network administrators of a large financial institute. He was primarily interested in the patterns of communication across group and organization boundaries. As subsequent studies also observed, his work illustrates the complexities of network (infrastructure) management work, pointing to the challenges of coordination across various dimensions (time, space, and organizational). Maglio, Harber, and colleagues later carried out an extensive series of fieldwork (for over 4 years) visiting a number of different IT service delivery organizations (from web hosting to operating system support to computer security to data center operations) (Haber and Bailey 2007; Kandogan et al. 2012). Kandogan and colleagues offer an in-depth and thorough account of IT service delivery in major IT outsourcing companies worldwide. Based on series of observations and contextual interviews, they present a series of episodes about sysadmins everyday practices, showing a great deal of deal on the ways in which they go about planning, deploying, monitoring, and troubleshooting IT infrastructures as inherently collective practices. Blomberg (Blomberg 2011) investigated the effect of process changes in organizational transformation, which involved both synoptic and performative accounts of day-to-day transformation of a large IT service organization. And, de Souza et al. (Souza et al. 2011a; Souza et al. 2011b) investigated the coordination practices in service delivery of a global IT service organization.

IT infrastructures are a key component, if not the most important one, of service delivery. But, as traditionally, they are usually in the background—invisible and frequently taken for granted (Star and Ruhleder 1994). In particular, these service delivery studies take infrastructure as the underlying substrate upon which sysadmins operate. As such, it conflates the systems (be it, computers, norms,

## RENEWING SERVICES

protocols, and the like) that support and enable their work and those upon which ‘they’ service. But, as will be described next, this takes no notice of the dual and relational nature of infrastructure where it is concomitantly in the background and the focal point, standardized and customized; local and global; internal and external; technological and social; transparent and opaque. Studying IT service delivery as part of large infrastructure deployment and maintenance thus becomes critical.

The traditional notion of infrastructure often invokes the image of the underlying substrates that support human activities, such as, water/sewer pipe, roads, rail-tracks, bridges, computer networks, and the like. This paper nonetheless draws on Star and colleagues’ longstanding tradition on ethnography of infrastructure. Drawing on a number of ethnographic studies of what Star came to call, the studies of “boring things,” they explored the embedded, messy, and relational nature of everyday infrastructures. From the start, they were particularly interested in studying unusual, often considered dull topics, which had attracted little attention of ethnographers. Nonetheless, as they delved into the various ‘boring’ topics they start to uncover a myriad of critical socio-political-cultural practices that had been left untouched by the simple fact that they were deemed as less exciting, exotic, and the like. In doing so, they start to investigate and deconstruct what they called infrastructural barriers, namely, social-political decisions that could enable, hinder, reinforce, or destabilize (depending on one’s point of view) actions, processes, institutions, or even knowledge (Dourish and Bell 2007; Star 2002). By its own nature, infrastructures are thus biased toward orthodox, conventional, standardized, conservative ways of doing things. This led to the understanding that infrastructure is in fact a mean of constraining, building, and preserving knowledge—based on whatever standards or cultural values. In other words, infrastructure embeds specific institutional arrangements that ultimately reify their existence. Significantly, Star and colleagues take infrastructure as fundamentally a relational concept (Star and Ruhleder 1996)—it can only be comprehended in relation to a particular organized practice or perspective. Hence, it has no fixed meaning. “One person’s infrastructure is another’s topic, or difficulty” (Star 1999).

### A CASE OF SERVICE DELIVERY IN A SERVICE FACTORY

Big Service Factory<sup>1</sup> (BSF) is an IT service factory, a large-scale IT service operation company where hundreds or even thousands of support personnel manage complex IT systems comprised of thousands of (heterogeneous) servers, routers, and other IT equipment, often from multiple customers, concurrently (Souza et al. 2011b). BSF employees are very specialized, often focused on specific tasks repeated exactly many times a day. In particular, many of these professionals are sysadmins, who are responsible for some of the most critical functions necessary to maintain the customers’ IT infrastructures running well and efficiently. Despite the specialization, the work of sysadmins is highly dynamic, collaborative, and interdependent (Haber et al. 2011).

In the past two years, we have been investigating the everyday service delivery practices in BSF (Souza et al. 2011a; Souza et al. 2011b). We started our research right after the BSF implemented a new organizational model (see next section) aiming at achieving economies of scale by means of shifting team organization from competence-based rather than customer-based. Our research thus took place

---

<sup>1</sup> This is a pseudo-name for the actual company’s name

## RENEWING SERVICES

in the midst of this major “infrastructuring” (Hughes 1983) process whereby previously invisible (uninteresting) elements came to the fore while others faded into the background.

Methodologically, we implemented a dual research program, inspired by ethno-mining (Anderson et al. 2009), which integrated both qualitative (ethnography) and quantitative (data-mining) methodologies. For the quantitative representation, we devised an analytical tool, named Workload Profile Chart (WPC), which characterizes the overall performance of a team by plotting on a log-log chart the normalized values of assignment and resolution times respectively of every incident in a given period of time. With regard to the qualitative study, we carried out a number of meetings in which we presented the WPCs to participants (ranging from sysadmins to management) and served as mediators of the conversations between researchers and participants—it helped translate a host of odd operations (according to response-time depicted on WPC) into SLA negotiation practices. This approach thus allowed us to unpack both synoptic and performative accounts of service delivery practices. Blomberg has proposed this methodological approach for studying the transformation of a large-scale IT service organization (Blomberg 2011), which intertwines synoptic and performative accounts of transformation with a service organization. In so doing, we were able to trace the trajectories of changes over time while at the same time unpacking the ways in which changes are in fact enacted in people’s everyday actions. Our ultimate goal was to uncover the ways in which service practices were transformed in people’s day-to-day work

## THE BSF’S ORGANIZATIONAL STRUCTURE


The BSF organization used to be organized in departments focused on particular customers, i.e., a group of sysadmins of different platforms would provide support to only a particular customer. In this context, they were more likely to build an in-depth understanding of important IT components of the customer to which they deliver services. On the other hand, the overall organization was unable to achieve economies of scale by sharing resources among customers, and second, technical knowledge was often not shared among those groups. Due to the highly competitive market, about seven years ago the BSF international started to implement a new organizational structure based primarily on technical competencies: namely, departments based on common skills, competencies, and activities performed. For example: UNIX group, responsible for dealing with issues with UNIX-based server systems; and Security group, responsible for accountability, security updates and similar issues across operating systems. These departments (or service pools) are responsible for handling incidents, where an incident is defined as “any event that is not part of the standard operation of a service and causes, or may cause, an interruption to or a reduction in the quality of that service” (The ITIL Open Guide, 2011). Each incident (or service request) is managed through an IT object often referred to as a ticket, which is a record that aggregates all the key information about the incident. Tickets describe an interruption (or reduction in service quality) of a particular service offered by the BSF to its customers formally agreed upon via legal contracts. Within the BSF, there is also a department responsible for bringing in new customers or selling additional services to existing ones. One of these teams is the customer account team whose job is to manage the overall relationship with the customer. They interact with both BSF employees and customers negotiating prices, contracts, and monitoring the quality of the services delivered.

Establishing well-defined roles is another aspect of the process standardization. For instance, to concentrate the information and knowledge specific to particular customers, the model proposes the

## RENEWING SERVICES

role of the customer experts (CEs), responsible for understanding in more detail the IT environment and the organizational structure of particular customers. However, given the large number of customers, some BSF employees are CEs for up to 5 different customers. In addition, within each department, sysadmins are grouped into three different job categories according to their level of expertise required to carry out the work. Another role is of the pool dispatcher, in charge of ticket assignments and monitoring; through the latter seems to be a collective responsibility, shared with sysadmins, duty-managers, and the like. Their assignment is based on their nature and difficulty level (or complexity) matched with a sysadmin's resolution capability (i.e., expertise). Figure 1 illustrates the process of ticket dispatching. Typically, tickets are created by customer requests and described on the IMS by help-desk personnel or generated by monitoring alerts (automated scripts) and routed to pools through the IMS. Then, based on the dispatching strategy adopted by the dispatcher working in the service pool, every incoming ticket is moved from the incoming queue to a sysadmin's who is capable of solving the incident appropriately.

The dispatcher thus plays a central role in service delivery, for her primary function is to match incident resolution requirements (based on the nature and difficulty of the problem at hand) and the best, available resource to tackle it, in a timely fashion. One of their main concerns is to comply with the SLA associated with every ticket. In so doing, they set the cadency of work performed by each pool. They set the *assignment time* of incidents. In juxtaposition lies the time it takes for a sysadmin to actually solve the incident—the *resolution time*. Incidents are recorded and represented by the incident management systems as *tickets*. Hence, tickets can be thought of as service requests. Considering that a ticket, say, with high severity might have the SLA of 30 (or even 15) minutes, any exceeding time spent *assigning* a particular ticket impacts its *resolution* time, thereby putting pressure on sysadmins. As will be discussed next, this is a major point of contention amongst different groups as they negotiate who fault it might be for missed SLAs.



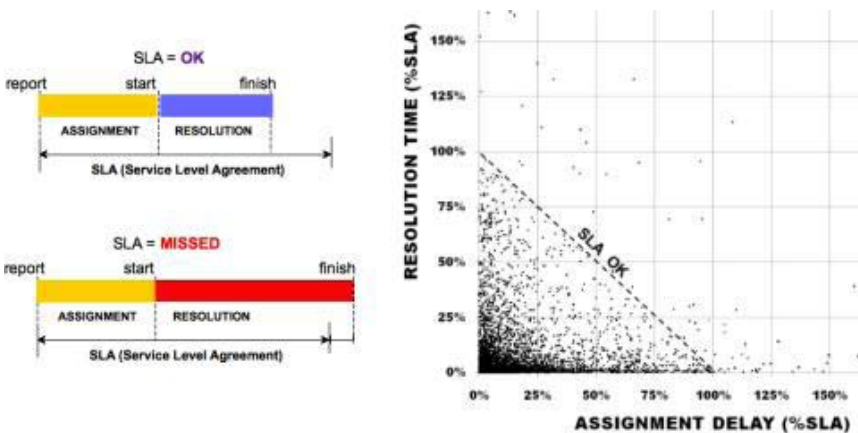
**FIGURE 1.** Ticket dispatching process.

## RENEWING SERVICES

Throughout the entire IT service delivery processes, a major concern is maintain IT infrastructures running uninterrupted and ready-at-hand, and whenever necessary to re-establish normal operations as *quickly* as possible with the least possible impact on customers' business or their clients. Hence, response time is a key metric of service quality. Contracts between IT service factories and their customers establish formal instruments to regulate the 'quality' of the service to be provided, clearly determining how fast incidents have to be investigated and resolved to guarantee the least possible disruption in IT infrastructure. To this end, SLA is a formal contract between parties that establish the maximum time within which a certain type of incident must be addressed and resolved. Typically, a formal classification of incidents is defined to associate types of incidents with their particular impact on a customer's organization due to service unavailability. Failing to abide to the SLAs of high priority tickets may result on fines and penalties.

## THE X-RAY OF SERVICE POOLS

Aiming at investigating ticket dispatching, we gathered from the IMS database tickets of four main service pools equivalent of approximately 10 months of work. In attempting to create a more effective visualization of the BSPF's overall 'performance' relative to ticket *assignment* and *resolution*, we first devised a scatter plot representation, where each point (x,y) represented a percentage of SLA spent during *assignment* and *resolution* times, respectively. In so doing, we were able to draw a first picture of the ticket distribution (see Figure 2b). First, we observed a large number of tickets at the lower part of the chart; that is, tickets being assigned and resolved under 25% of SLA (*assignment* and *resolution*) times. We needed however to come up with a better way to visualize these tickets.



**FIGURE 2. (a) Ticket assignment and resolution representation; (b) Ticket assignment and resolution chart.**

We shifted from a scattered plot representation to a log-log chart (see Figure 3). This new visualization helped us reduce the clutter on the chart by magnifying the 'distances' for low order

## RENEWING SERVICES


*assignment* and *resolution* times. To further improve our reading of the chart—in particular, we were interested in answering the question of “how to visualize the number of tickets occurring in similar (x,y) location”—we decided to treat our chart as a grid. As such, each cell (or bean) of the grid contained the sum of tickets within delineated quadrant. It then turned into a heat map representation. We called that new chart, Workload Profile Chart (WPC), but described it as the “X-Ray” of service pools. In practical terms, it can be thought of as a characterization of the performance of service pools relative to the overall *assignment* and *resolution* times with respect to a customers’ SLA. Hence, the allusion to an “X-Ray,” It is a picture of the “guts” of a service-line, namely, a representation of its productivity and service quality that allows, for example, to map transformation impacts and to *diagnose* service delivery problems. In short, the WPC depicts of the workload profile density matrix with a grey scale associated to different ranges of ticket concentration.

## UNDERSTANDING PRACTICES

The field study thus examines in greater detail particular ‘behaviors’ of these service pools and the ways in which everyday practices might have affected the manifest characteristics observed in their particular WPCs (Figure 3). It thus enabled us to gain critical insights relative to the distinct processes, practices, and problems that shape the ways in which different pools handle tickets. We believe that these insights help illustrate the factors and circumstances that impacted those service pools over the period of this research, which in turn resulted in particular ticket distribution manifested in the WPCs. It’s noteworthy that we did not only look into the overall aggregate of 10 months of work; we also examined the month-to-month transformations as well as differences across work-shifts. For brevity’s sake, we will only represent the overall aggregate of tickets handled from March to December 2011, but the interviews, conversations, and analyzes utilized the overall spectrum of representations at hand.

This paper is particularly interested in the hectic and extreme ‘behaviors’ that emerged from the inspection of WPCs, nonetheless. For instance, service pool 3 shows an odd behavior of that tickets tend to be resolved at the end of SLA (areas 3a), however easy to solve they might be. We found for instance an interesting practice of “holding” the ticket—for certain types of incidents, sysadmins would stop the time just before the SLA is met so as to have more time for figure out a solution. One can clearly see the number of interesting questions that this simple fact warrants. Why are sysadmins waiting until the last minute to answer these ‘easy-to-solve’ tickets? In contrast, service pool 4 seems to operate much more comfortably, where tickets sit around the “comfort zone.” The “comfort zone” refers to the areas found at the center of the WPCs [1%-10%, 1%-10%] where tickets are “comfortably” assigned and resolved (see the “squares” in Figure 3). A high concentration of tickets in the comfort zone seems to characterize a service pool that is working smoothly and comfortably. In fact, it warrants some interesting tensions/questions. Is this due to over-staffing? Or is this ‘behavior’ due to the nature of software application incidents? We asked, participants of the study demonstrated that the notion of a “comfort zone” is much more contingent on the everyday configurations of a particular service pool than a broad, static definition. But, at its core, we are particularly interested in investigating how a particular pool can be best staffed or structured to handle incoming incidents. Hence, we will describe briefly a few observations from the WPCs and the stories we were told and explanations given relative to these particular characteristics such as to illustrate how we moved back and forth from the WPCs to the field.

## RENEWING SERVICES



**FIGURE 3:** The X-Rays of four service-pool from March thru December 2011 (areas of interest highlighted)

**Coping with structural changes** – When analyzing and discussing with participants the month-to-month changes in WPCs, we were able to capture relevant insights relative to major organization transformations that took place throughout the last 10 months of 2011. In analyzing the changes in the distribution of tickets over time for service pools 1 and 2, we came with the hypothesis that some major events contributed to two major ticket dispersions. As soon as we showed research participants the time evolution diagrams of the pools, they brought up that in mid-2011 a new, large customer outsourced its IT infrastructure to the BSF. The sudden demand for IT resources (storage, servers, infrastructure, and the like) momentarily affected the overall operations of the BSF. Such a demand contributed to the generation of large number of tickets. They became overwhelmed with system requests. The overall service delivery system became provisionally unstable (i.e., the observed ticket dispersion), as the BSF run to increase its infrastructure capacities. After a few months, it is noticeable



## RENEWING SERVICES

that the system started to find its new equilibrium as new IT resources are brought in to meet the demand. Service pool 3 was, on the other hand, affected by the lack of skilled workers. Research participants pointed out the difficulties of finding skilled workers for the particular backup & restore systems in use in the BSF. By the end of the year, the service pool was able to hire new employees but mostly junior—without prior system skills.

**Mismatched Standards** – Service pool 3 manifests a particularly distinct ‘behavior’ as compared to the others, namely, a large number of tickets handled solely at (and around) the limit of the SLA. As sysadmins pointed out, backup & restore are long-lived processes which reflect the fact that it takes time to carry out these tasks. In fact, at times, it might take longer to carry out the resolution task (such as, restore a backup) than the SLA time limit permits. Also, they are usually scheduled for late-night shifts. Hence, there is a high degree of uncertainty as to the ability of the pool to respond on schedule according to a particular SLA. To workaround this problem, sysadmins informally agreed internally to “stop the clock,” (or as they put it, “to hold SLA.”) That is, they stop the SLA time ticking while waiting for the operation to take place.

This demonstrates a mismatch between the SLA agreed upon between the BSF and its customers and the actual technological realities, for instance, customer’s slow backup devices. Interestingly, as for the tickets orbiting the region about the comfort zone square, in the interview, sysadmins of service pool 3 pointed out that these are more likely to be the tickets correctly timed whereas the ones at the lower regions more likely represent those that clearly were not adequately timed. This resulted in a conversation about the effectiveness of SLA metrics. They pointed out that, in their case, SLA should not account simply for the percentage of incidents responded in time, but the overall availability or unavailability of the systems in questions. That is, they pointed out to the problem that their service pool might miss the target SLA (e.g. 5% of the incidents), although the actual impact of such a miss might represent just 1% impact on the system overall availability. As they put, it would not matter whether 2 or 20 hours has lapsed, but the aggregate number of incidents that missed the target SLA. Significantly, it affects not only the overall measure of performance, but also the choices sysadmins make relative to what incident to respond next. Hence, the behavior we observed manifests their choices of tackling the incidents that are about to ‘expire’ as opposed to focusing on those with potentially greater impact on the customer.

**Not my fault, ticket exchange practices** – In general, we observed a large number of tickets being assigned at last minute, where some were quickly solved (and sometimes dismissed, as we came to learn) and some miss the SLA (in particular, see 1a and 3a in Figure 4). That is, these are tickets that are taking more time to be assigned than resolved (or simply taking too long to start being resolved). For one, the process of ticket assignments (or dispatching, as previously described) heavily relies on the dispatcher’s capacity of interpreting the problem described in the ticket—the nature of the incident—in order to efficiently and correctly assign it to the best available sysadmin in the pool to tackle it. But, at times, tickets’ assignments get delayed resulting in reduced time-window for their resolution. Given that the SLA accounts for the total time (*assignment* plus *resolution* times), the sysadmin who is assigned to a belated ticket will have less time and more pressure to get the work done.

For example, a ticket might be mistakenly placed in a pool’s queue by the help-desk operator who first answered the customer’s request. (In fact, we learned that the help-desk personnel tend to place incidents in the OS queues whenever they are unsure of the actual nature of the incident at hand.) In fact, it might sit on the dispatcher’s or sysadmin’s queue for a while until they realized that it was on

## RENEWING SERVICES

the wrong queue. It thus has to be re-routed to the appropriate queue in order to be dispatched once again, while the time is still ticking. A ticket can actually go back and forth between different service pools a couple of times before it gets resolved. When the new dispatcher or sysadmin finally receives this belated assigned ticket and is unable to resolve it in time, he will most likely contest it so that his pool does not take the blame for the missed ticket, given that the “offender” is from elsewhere—so they say.

At times, a dispatcher or a sysadmin might need more information in order to resolve the incident than what was originally described in the ticket. She might then consult other colleagues or ask the customer for more specific information. Again, the time continued ticking. In case this was perceived as having hindered the resolution process and prevented her from meeting the SLA, she might ask to “expurgate” this missed ticket from the DB so that it would not count toward the pool’s monthly allowance. Missing the SLA monthly target will incur in financial penalties and most likely to affect customer’s relationship. At the service pool as well as organization levels, monitoring the SLA status of individual customers is an ongoing concern so that attention can be paid to those that are at risk of missing the SLA target at the end of the month. Customers’ accounts at risk become the focus—more resources are allocated, incidents are prioritized, account business managers scrutinize pools’ actions, and more—so that no more incidents are lost and the service response performance improved. In this context, the possibility of purging (or expunging) missed tickets becomes a powerful mechanism of coping.

**Automated systems and the flood of false alarms** – While our study focused on the dynamics of everyday practices as affected by the nature of incidents as well as organization structures, we also looked into the affect of the tools sysadmins use to carry out their work. For instance, service pools 1 and 2 (Figure 4) respond to OS related incidents. We can observe quite similar ticket distribution characteristics including a large number of false alarms or easy-resolution tickets (the large concentration of tickets in areas 1A and 2A in the figure). Significantly, OS experts pointed out that a great number of tickets are automatically generated via OS scripts, which monitor various states of the servers. Hence, the observation that this is due to false alarms and tickets of easy resolution. One of the interviewees, for example, pointed out that such a large number of easy-resolution tickets is in fact a typical characteristic of OS service pools in which CPU usage, HD space, and the like are constantly monitored by automated scripts. Interestingly, as transpired throughout our studies, this is a point of contention—a clear conflict between a team’s resolution goals and its performance metrics. As team performance is in part measured by meeting a target SLA (percent of tickets resolved in time), to reduce the ‘number’ of tickets that can be easily resolved directly affects (percentagewise) their performance, as they will be left with a larger number of difficult tickets to resolve. A clear distortion created by this simple, inconspicuous performance assessment.

Similarly, as discussed earlier, backups (service pool 3 in Figure 4) are long-lived operations—it simply takes time to backup and restore a machine. Hence, they might be subject to interference and disruption from other unscheduled or unanticipated system operations, such as, OS updates, reboots, and the like. These faulty operations then generate a large number of tickets as backups are interrupted in the middle of their processes. Provided that backup is a recurring operation (daily, weekly, and monthly), it creates a cascading effect as automatic tickets due to faulty backup pile on top of next cycle’s scheduled backups.

## RENEWING SERVICES

Participants of the fieldwork (i.e. sysadmins) pointed out that this large number could also represent that they might need to re-calibrate their metrics and monitoring scrip-rules since these were too strict. In this respect, the WPC quickly became a tool whereby sysadmins were able to reflect on the characteristics and behaviors of their particular pool and daily practices. The high volume of tickets on the ‘edge’ was also a topic of discussion by OS sysadmins. A lack of personnel (number and skills) was identified as the possible cause. In one pool, the limited number of people with adequate skills was pointed out as a main concern. Staffing quickly becomes an ‘impromptu’ topic of discussion in the interviews, meetings, and informal conversations. As technologies develop, for instance, the development of virtual machines (VMs), cloud computing, and the like, it becomes increasingly ‘easy’ to ‘clone’ machines, creating literally at the touch of a button new technological infrastructures that can quickly address business and system demands and opportunities. But, on one hand, these technologies are not resource limitless (i.e., CPU, memory, HD, and the like have limits), and, on the other, they demand more human attention (i.e., people capable of keeping these systems running, as previously described). Hence, the emergence of these new technologies that enables the business to quickly (more often, hastily) bring in new clients falls short of creating socio-technical conditions for their actual assimilation and integration.

## FINAL REMARKS ON HETEROGENEITY

The integration of the WPC inspection method with fieldwork has helped unveil some of the actual meanings and practices of IT service delivery that were manifest characteristics of the various configurations of the WPCs. We were able to observe how organizational changes (process standardization, new clients, re-orgs, operational optimizations, and the like) affected the ways in which those involved (e.g. system administrators, customer experts, infrastructure architects, dispatchers, and the like) were able (or unable) to accomplish their work. On the other hand, WPCs revealed a number of unarticulated practices that would have gone unnoticed in our investigations for the simple fact that they are unaccounted for and not aligned with established process models and standard procedures. All in all, we believe that the WPCs allowed us to view some of the *traces* left behind by the dynamic and complex interactions among this network of actors (animated and otherwise) that constitute the everyday service delivery practices.

Hence, the method has unveiled part of the heterogeneity of service delivery that rests not only on the configuration of the work settings—by means of its complex and heterogeneous arrangements of computer systems, network cables and routers, software applications, and more—but also on the very nature of the work practices—by means of people interdependences, organizational structures and standards, and socio-technical infrastructures. In light of Latour & Law’s notion of heterogeneity, we can exam service delivery as complex assemblages of heterogeneous materials. The ‘social’ is not an aspect of the work processes; instead it emerges as *traces* from the interactions among those materials, according to Latour (Latour 2005). However, this should be considered not as a departure from Star et al.’s theoretical and methodological approach for studying ‘infrastructures,’ rather the notion of heterogeneity is an analytics lens that allows us to draw together this vast set of actors (animated or otherwise) to play important roles in determining the usefulness of IT service delivery systems. From the more complex, abstract server operations to client’s business critical applications to monthly performance metrics to the more mundane (often tedious) process norms, the everyday routines in a

## RENEWING SERVICES

service delivery organization is comprised of more than simply keeping the ‘infrastructure’ running, but managing and fixing the relationships among its various components.

By introducing heterogeneity in the study of IT service infrastructures we aimed at contributing to the long line of ethnography of infrastructure research. We were able to account for the complex assemblages of heterogeneous materials that come into play in an attempt to move beyond the dichotomy between technology and social and show the ways in which these materials affect ongoing everyday accomplishments. Methodologically, we showed the critical synergy between quantitative and qualitative methods in support of large-scale work practice research.

## REFERENCES CITED

- Anderson, Ken, Dawn Nafus, Tye Rattenbury, and Ryan Aipperspach  
2009 *Numbers Have Qualities Too: Experiences with Ethno-Mining*. Ethnographic Praxis in Industry Conference Proceedings 2009(1):123-140.
- Blomberg, Jeanette  
2011 *Trajectories of Change in Global Enterprise Transformation*. Ethnographic Praxis in Industry Conference Proceedings 2011(1).
- Bowker, Geoffrey  
1994 *Information Mythology and Infrastructure*. In *Information Acumen: The Understanding and Use of Knowledge in Modern Business*. L. Bud-Friedman, ed. Pp. 231-247. London, UK: Routledge.
- Bowker, Geoffrey C., Karen Baker, Florence Millerand, and David Ribes  
2010 *Toward Information Infrastructure Studies: Ways of Knowing in a Networked Environment International Handbook of Internet Research*.97-117.
- Dourish, Paul, and Genevieve Bell  
2007 *The infrastructure of experience and the experience of infrastructure: Meaning and structure in everyday encounters with space*. *Environment and Planning B: Planning and Design* 34(3):414-430.
- Haber, Eben M., and John Bailey  
2007 *Design guidelines for system administration tools developed through ethnographic field studies*. In Proceedings of the 2007 symposium on Computer human interaction for the management of information technology. Pp. 1. Cambridge, Massachusetts: ACM.
- Haber, Eben M., Eser Kandogan, and Paul P. Maglio  
2011 *Collaboration in system administration*. *Commun. ACM* 54(1):46-53.
- Hughes, Thomas Parker  
1983 *Networks of Power: Electrification in Western Society, 1880–1930*. Baltimore, MD: Johns Hopkins University Press.
- Kandogan, Eser, Paul Maglio, Eben Haber, and John Bailey

- 2012 *Taming Information Technology: Lessons from Studies of System Administrators*. Oxford University Press.
- Latour, Bruno  
2005 *Reassembling the Social: An Introduction to Actor-Network-Theory*. Oxford, UK: Oxford University Press.
- Law, John  
1992 *Notes on the theory of the actor-network: Ordering, strategy, and heterogeneity*. *Systemic Practice and Action Research* 5(4):379-393.
- Sandusky, Robert J.  
1997 *Infrastructure management as cooperative work: implications for systems design*. In Proceedings of the international ACM SIGGROUP conference on Supporting group work: the integration challenge. Pp. 91-100. Phoenix, Arizona, United States: ACM.
- Souza, Cleidson R. B. De, Claudio S. Pinhanez, and Victor Cavalcante  
2011a *Knowledge and information and needs of system administrators in IT service factories*. In Proceedings of the 10th Brazilian Symposium on on Human Factors in Computing Systems and the 5th Latin American Conference on Human-Computer Interaction. Pp. 81-90. Porto de Galinhas, Pernambuco, Brazil: Brazilian Computer Society.
- Souza, Cleidson R. B. de, Claudio S. Pinhanez, and Victor F. Cavalcante  
2011b *Information needs of system administrators in information technology service factories*. In Proceedings of the 5th ACM Symposium on Computer Human Interaction for Management of Information Technology. Pp. 1-10. Cambridge, Massachusetts: ACM.
- Star, Susan Leigh  
1999 *The Ethnography of Infrastructure*. *American Behavioral Scientist* 43(3):377-391.
- 2002 *Infrastructure and ethnographic practice: Working on the fringes*. *Scandinavian Journal of Information Systems* 14(2):107-122.
- Star, Susan Leigh, and Karen Ruhleder  
1994 *Steps Towards an Ecology of Infrastructure*. Computer Supported Cooperative Work, Chappel Hill, NC, USA, 1994, Pp. 253-264.
- 1996 *Steps Toward an Ecology of Infrastructure: Design and Access for Large Information Spaces*. *Information Systems Research* 7(1):111-134.