Modeling Service-Level Requirements: a Constancy Perspective

Gil Regev[1, 2], Olivier Hayard[2], Donald C. Gause[3], Alain Wegmann[1]
[1]*Ecole Polytechnique Fédérale de Lausanne (EPFL), School of Computer and Communication Sciences, CH-1015 Lausanne, Switzerland, {gil.regev, alain.wegmann}@epfl.ch*
[2]*Itecor, Av. Paul Cérésole 24, cp 568, CH-1800 Vevey 1, Switzerland*
{g.regev, o.hayard}@itecor.com
[3]*State University of New York at Binghamton and Savile Row, LLC, U.S.A*
dgause@stny.rr.com

Abstract

IT service requirements offer a seemingly classic Requirements Engineering (RE) problem. But, when attempting to solve it with RE methods, we are faced with difficulties. RE methods encourage us to identify the functional and non-functional requirements of a service. Industrial service-management frameworks, however, use a different vocabulary. ITIL, one of the most prominent service-management frameworks, refers to service utilities and service warranties. In this paper, we propose a method for modeling warranties as a function of the service constancy expected by stakeholders and the threats to this constancy. We identify four kinds of warranties: express, implied, tacit and pending. We thereby seek to bridge the gap between service-management frameworks and RE methods and to improve the practice of service management in organizations.

1. Introduction

In our consulting businesses, we see that the service paradigm provides a major opportunity to introduce Requirements Engineering (RE) practices in organizations. Indeed, service initiatives are often spearheaded by IT departments as a way to improve the overall workings of an organization and its use of information technology. For this to happen, a mutual understanding of business needs, constraints and the potential of IT has to be created between IT departments and their stakeholders. This is a classic RE preoccupation, but existing RE methods are not well positioned to be used in these endeavors. We have recently uncovered that one of the reasons for the difficulties is that IT service frameworks do not use the methods and concepts developed in the RE community; they rather depend on ad-hoc specially developed methods and concepts. As a result, even though the opportunity for augmenting the awareness of RE in organizations exists, the actual use of RE methods may not be increasing.

In this paper, we concentrate on the concepts of service utilities and service warranties as they are defined in the IT Infrastructure Library (ITIL) [6, 11], one of the most prominent IT service frameworks of the day. As we argue elsewhere [8], these concepts are similar to the concepts of Functional and Non-Functional Requirements (FR and NFR). There are, however, subtle differences between these concepts. The use of the term warranty, for instance, has the merit of attracting the attention to the guarantees that stakeholders expect from a service provider in a way that the more general term NFR does not. Moreover, ITIL has a whole section devoted to RE, which contains an explanation of FR and NFR but that does not attempt to link them with Utilities and Warranties.

Earlier we showed [8] that service warranties can be understood as the assurance that a service will adhere to the norms (stable states) that stakeholders are accustomed to, (i.e. the stakeholders' tolerance to gaps between their expectations and the actual service they receive). The underlying assumption is that stakeholders expect a large amount of constancy[1] from the services they use. Now, we focus more specifically on the identification of these norms and tolerances, and on the alignment of IT service warranties with the business service warranties and their evolution.

We sought to validate our thinking with an example where there are clear (and not so clear) warranties and where norms and tolerances of stakeholders can be openly investigated. The Swiss federal postal service presented a readily analyzable and comprehensible example that is neither too complicated nor too simple. We made three visits to the Swiss Post sorting center that serves the western part of Switzerland. The visits

[1] Constancy is a synonym of stability and we use both terms interchangeably. We do favor constancy in a more general discussion because it implies less rigidity.

Presented at the 17th IEEE International Requirements Engineering Conference (RE'09), Atlanta, GA, August 31 - September 4, 2009

included an interview with the center's manager and tours of the center's facilities where we were able to observe the actual work being done and to interact with the center's employees. We present our interpretations of the reality we observed and how they can be modeled with our proposed framework. We limit our discussion to the availability warranty.

In Section 2 we provide the theoretical background of system constancy and its relation to warranties. In Section 3 we propose four kinds of warranties. In Section 4 we illustrate our framework with the postal service example and explain the warranties we identified. In Section 5 we link our work with previous research. In Section 6 we summarize our work.

2. Systems, Warranties and Constancy

Systems Theory [9] holds that a system is an island of order within a sea of disorder. Survival is therefore defined as the ability for a given observer to maintain this order. Closed systems, by virtue of the second law of thermodynamics are said to evolve over time toward disorder (an increase of entropy). In order to survive (maintain negative entropy), a system must have relationships with other systems. A system draws energy, information and matter from the systems with which it has relationships in order to maintain its internal order (this is the concept of open system).

The concept of open system implies that systems must accept input from other systems in order to survive. In doing so, a system becomes dependent on the stability of the input it receives from the other systems' output. For a network of interdependent systems to function without failure, each system must control the input it receives from the other systems and the output it provides to them so that both input and output remain within a limited region in the system's descriptive state space, called the homeostatic region. Homeostasis, the act of maintaining residency within this homeostatic region, is referred to as the maintenance of identity [7, 13].

We refer to the homeostatic region as a norm (a state that changes slowly over time [7]) and the tolerance of the system to variation of this norm. The maintenance of a system's output within the tolerances associated with the norm is the warranty given by the system to the systems that depend on it.

From the viewpoint of service designers, the goal is to design service warranties that fit the norms and tolerances of its stakeholders. The world of design can be seen as consisting of context and form, where the form is the product of the designer's activity and context is everything else [2]. Hence, service designers must design the service warranties (the form) so that they fit the context, the norms and tolerances of stakeholders.

We model this design using the Systemic Enterprise Architecture Methodology (SEAM) [10]. SEAM enables us to create an enterprise model that represents a company's environment, its roles, its organization, its IT systems functionality and their construction. The SEAM notation is inspired by UML.

Some of the SEAM modeling elements are visible in Figure 1. The arrow box represents a system in the general sense, i.e. a set of inter-related elements. We model two kinds of hierarchies, the organizational hierarchy (which system belongs to which outer system) and the functional hierarchy (which action belongs to which global action). In SEAM, the organizational hierarchy is shown when a system is represented as a black box (a system as a whole) and as a white box (system as a composite). We use the representation of a system as a whole to analyze the service offered by an organization to its stakeholders, and the system as a composite to analyze the way the organization implements the service.

3. Understanding Service Warranties

The Cornell University Law School Legal Information Institute (LII) provides legal definitions for the two readily accepted warranty types, the *express* or *written warranty* and the *implied warranty* [1]. For the purposes of this paper, and in order to put these terms into a systems design context, we relax the legal formalities while paraphrasing the legal definitions and adding two additional warranty classes, *tacit* and *pending*. In the following list we describe the resulting four service warranties and link them with the four classes of requirements defined in [2], the obvious, subtle, invisible and unknowable.

Express Warranty – Is a promise or commitment, formally stated, made by the seller (system designer in our case) to the buyer (system users). This class of warranty describes the *obvious* context relevancies of the system to be designed. The obvious contains critical information that is widely known and understood by designers and users alike. This would appear to be the safe class and it is relative to the other three classes. The risk we run with this class is associated with dropping important features and functions because of conflicting priorities rather than lack of awareness.

Implied Warranty – Is an informal understanding by designers and users based on existing norms, their respective needs and past experiences, as well as service expectations based on marketing communications. These warranties identify the critical context relevancies that are available only after we have applied special processes and disciplines to obtain them. They are inherently *subtle*. The subtle contain critical relevancies that are known, and even obvious, to various user constituencies but not recognized by the designers. The expectations of some of the user constituencies will not be met, thus leading to disappointment and risk of service failure. At least one source of difficulty we have observed in business systems and services occurs in the design and development of strategic business services. Business requirements are developed and passed on from system analysts to designers by way of system requirements. But the designers may lose the coupling between the business needs and the end user needs as they begin to implement more detailed functional requirements.

Tacit Warranty – We can think of this as a warranty that has not been recognized and is therefore not consciously incorporated into the design process, it is *invisible*. The invisible are the relevancies that deal with environmental issues that are available to us but only after extreme, or at least, extra effort.

Pending Warranty – The final warranty class contains the critical system design relevancies that are not likely to surface until the system is released and put into operation. They are in a class of unknowable context relevancies. The unknowable, by definition, would appear to be next to impossible to discover until the service is put into operation. After all, this is the class that deals with surprising and often expensive unintended consequences.

We add these two phantom warranty categories, tacit and pending, in order to raise our awareness of the existing design uncertainties and risks. Moreover, in spite of their seemingly hopeless names, tacit (invisible) and pending (unknowable), we can actually take measures to mitigate the risk and even eliminate some of their uncertainties. The more a service is innovative, the more we can expect tacit and pending warranties; simply because innovative services, by definition, are outside the homeostatic region with which stakeholders are familiar and for which express and implied warranties exist already.

The ultimate goal of defining these warranty classes is to recognize as many critical design issues as early as possible and to move, wherever possible, from the unknowable to the invisible, and up through subtle to the obvious.

4. The Swiss Post Example

Swiss Post[2] runs the federal postal service within the borders of Switzerland. It is responsible for letter and parcel delivery to businesses and private people. The Swiss Post sorting center, which we analyze in this example, processes about 3.5 million letters a day. In this analysis we aim to model the warranties expected by the Swiss Post stakeholders and to explain how these warranties are supported by the warranties that the IT service of the sorting center maintains. We analyze (1) the service provided by Swiss Post to its customers, (2) the service provided by the sorting center to Swiss Post's internal stakeholders, and (3) the service provided by the sorting center's IT systems to the sorting center's internal stakeholders.

Swiss Post Service to its Stakeholders

First-class letters addressed to a P.O.Box must be delivered by 7:30 AM. This defines the availability of the service in terms of time and space.

Figure 1 represents Swiss Post in its market segment. Swiss Post is represented as a whole, focusing on the service it provides to its stakeholders (customers). Simplifying the extremely complex context of Swiss Post, we represent only two customers, Letter Sender and Letter Receiver. Warranties appear in the model as properties tagged with the "warranty" stereotype. In Figure 1, the warranties we consider for the First-class Mail Delivery Service are: *Pickup & delivery at any Swiss location* and *P.O.Box delivery at 7:30 AM next day*.

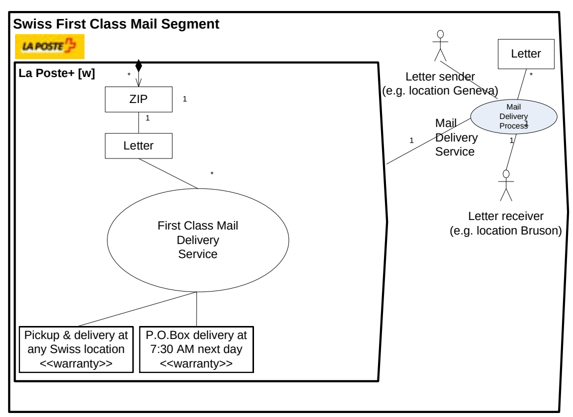

Figure 1 Swiss Post Service to its Stakeholders

There are numerous threats to the stringent guaranty of next-day delivery time for first-class letters. Transportation means (trucks and trains) can fail due to many causes such as mechanical failure, bad weather and traffic jams. Sorting errors by humans and machines can divert individual letters or complete shipments from their expected destination. Hardware or software failures can break sorting machines and

[2] http://www.poste.ch/en/index.htm

automated inventory systems, slowing down or halting the sorting process.

Sorting Center Service to Swiss Post

Figure 2 represents Swiss Post as a composite, a much simplified model of the Swiss Post internal organization.

The mail distribution process is the management of the circulation of shipping boxes. A shipping box is a plastic box containing about 300 sorted letters. The process involves three separate organizations, the receiving organization, the sorting center and the distribution organization.

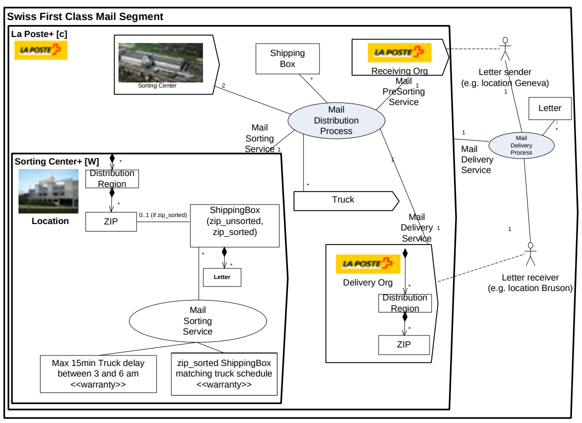

Figure 2 Sorting Center Service to Swiss Post

The receiving organization is in charge of aggregating individual letters into shipping boxes. It pre-sorts the mail according to its class (first, second, third) and separates mail that requires manual sorting. These responsibilities are not represented in our simplified model. Trucks transport the mail from the receiving organization to the sorting center.

The sorting center takes unsorted letters in shipping boxes, sorts the letters and places letters with the same destination into a shipping box that now has a destination in the form of a zip code. Shipping boxes are loaded into carts and the carts are loaded into trucks for shipment to their destination. The express warranties provided by the Sorting Center to the Swiss Post internal stakeholders are: *Max 15min truck delay between 3 and 6 AM* and *zip-sorted ShippingBox matching truck schedule*.

Trucks that transport mail follow a carefully planned schedule. Keeping to the tight schedule is paramount to upholding the first-class mail delivery warranty. If a letter is sent from Geneva to the alpine village of Bruson, for example, several mechanisms need to be interconnected for it to be in a P.O.Box at the Bruson post office by 7:30 AM the next business day. There are probably several incoming local collection trucks that need to connect with other trucks or trains so that the letter arrives at the sorting center by nightfall. On its way out, the letter has to be on a truck outbound to the distribution center in the city of Bex[3] where it has to be placed on a connecting truck for it to be in the Bruson post office before 7 AM and be deposited at its destination P.O.Box by 7:30. Delaying the truck for more than 15 minutes is very likely to prevent it from reaching Bex on time for its load to be transferred to the connecting trucks. Between 3 and 6 AM, the last trucks bound for a given region are scheduled to leave the sorting center. If our letter to Bruson misses the last truck it will be delayed by one day and the Swiss Post warranty to its customer will not be fulfilled (*P.O.Box delivery at 7:30 AM next day*, see Figure 1). This gives us one of the norms we were looking for: match truck schedule. The tolerance is a delay of no more than 15 minutes in this schedule from 3 to 6 AM. Hence the service warranties shown in Figure 2.

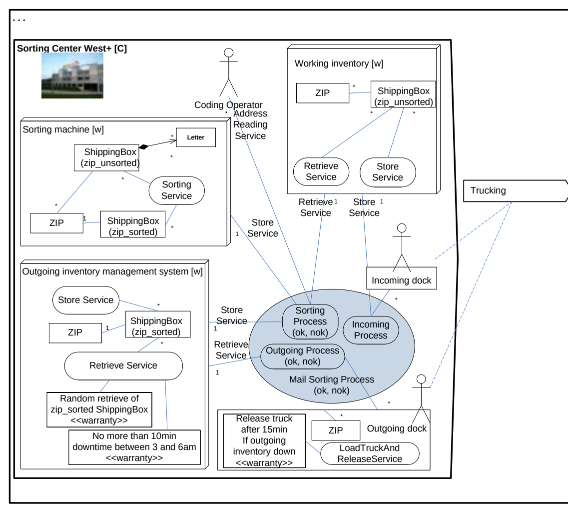

Figure 3 IT Service to Sorting Center

IT Service to Sorting Center

In Figure 3 we represent the sorting center as a composite, which enables us to show the service provided by the IT to the sorting center's stakeholders. The sorting process is composed of three activities, incoming, sorting and shipping. The incoming activity takes the pre-sorted shipping boxes from the incoming dock and places them in the working inventory.

A partially automated sorting activity outputs the letters sorted by ZIP code into shipping boxes that are

[3] http://en.wikipedia.org/wiki/Bex

placed in the outgoing inventory. During the shipping activity, shipping boxes with a shared destination are pulled out of the outgoing inventory, placed in carts with other boxes sharing the same destination and moved onto the outgoing dock. The sorting center's service warranty is limited to placing shipping boxes containing sorted letters in carts on the outgoing dock, ready to be loaded into trucks.

The outgoing inventory management system contains two warranties: *Random retrieve of zip-sorted Shipping Box* and *No more than 10min downtime between 3 and 6am*. The outgoing dock also has an implied warranty: *Release truck after 15min if outgoing inventory down*.

The movement of shipping boxes out of the outgoing inventory is done automatically by the *outgoing inventory management systems* (the sorting center has several such systems working in parallel). If, let's say, the outgoing inventory management systems fails to put the shipping box with our letter to Bruson on its corresponding cart by the time the last truck to Bex is scheduled to leave the dock, the sorting center's personnel has a dilemma. Should they delay the truck and wait for the shipping box (and maybe other similar shipping boxes) to be delivered to the dock or should they release the truck without the delayed shipping box? A tradeoff has to be made between delaying a few letters by one day versus delaying a truck load by more than 15 minutes which may result in a one-day delay for many letters in the truck. Hence, we have two warranties: (1) The outgoing inventory management system cannot be down for more than 10 minutes between 3 and 6 AM. (2) A truck is to be released after a 15 minutes delay, even if all of its planned content has not been loaded. The second warranty is a backup mechanism for the case when the warranty of the outgoing inventory management system does not hold. For Swiss Post customers, this is probably a tacit warranty, because they simply expect Swiss Post to distribute letters by mitigating the consequences of failures but are unaware what tradeoffs have to be made in the process.

The third warranty *Random retrieve of zip-sorted ShippingBox* corresponds to a different norm and tolerance. The outgoing dock is configured to serve a specific geographical region. For example, lines 1, 2 and 3 receive the mail to be sent to the region of Bex. When the sorting center began to operate, each outgoing inventory management system was configured to serve the lines that are physically nearest to its outlet so as to minimize the travel distance of the carts and to ship all the mail to the region in one batch. An unintended consequence was that when the outgoing inventory management system that served the region of Bex failed, the mail for a whole region was delayed. The tolerance for a whole region being denied of mail for a whole day is much lower than individual mail being delayed. The sorting center has changed its process and now distributes the mail randomly over the outgoing inventory management systems. This spreads a potential failure on multiple regions (which is far more acceptable). However, it comes at the expense of some inefficiency because the carts have to be ferried over longer distances.

This is then an example of a pending warranty. It was discovered only after the service was put into operation. Note that we have omitted the corresponding warranty in the model in Figure 1. Adding the warranty, *Minimize possibility that whole region is affected by a single outgoing inventory management system failure,* at the sorting center level is equivalent to making it an express warranty to the sorting center designers and a tacit warranty to Swiss Post customers.

5. Related Work

As we point out in the introduction, service warranties can be seen as a special case of NFRs. Our work, therefore, is closely related to the work of many RE researchers on NFRs, e.g. i*. i* has the advantage of modeling tradeoffs in terms of choice between softgoals, a very useful tool when, (as we have shown in the sorting center description), tradeoffs are built into the design of the service. i* models also show how actors depend on one another. We show this to be an important aspect of aligning warranties. Similarly, our work can help with the theoretical understanding of warranties and therefore improve the modeling of services with i*.

e3-value [3] is an RE method based on the analysis of value exchanges within a network of actors. e3-value is used for analyzing the value offered by a service and expected by stakeholders. Recent work concerning the alignment of business and software services can be found in [4].

Lauesen [5] briefly explores the issue of tacit requirements and the way they are handled in courts of law when customers do not have the impression of having obtained the product features they expected or could rightly expect. Lauesen proposes heuristics for domain obvious and domain non-obvious requirements analysis with respect to the complexity of the program that needs to be written.

Finally, SERVQUAL [15] is a framework for understanding customer expectations in order to improve service quality. SERVQUAL has been defined for marketing services.

6. Conclusions

The concepts of service utilities and warranties were recently introduced by ITIL. They provide a different view on requirements engineering for services than the traditional FR/NFR dichotomy prevalent in RE. By introducing these ITIL concepts to RE discourse, we seek to bridge the gap between service management as practiced in organization and RE theory.

Service warranties are commitments made by service designers to service stakeholders that the utilities the stakeholders have come to depend on will fit their tolerances for failure. In many ways warranties represent a small subset of the full system requirements relevant (in the eyes of the designer) to a potential stakeholder. They correspond to the requirements that the designer feels the stakeholder needs to understand based on the reasonable operation of the designed service. They also serve as disclaimers in a legal sense, thus establishing boundaries for which the designer accepts and denies responsibility.

The stakeholder, in turn, is faced with the reality of human expectations. A proper set of warranties must provide for consistency between designer requirements and stakeholder expectations. When we actively incorporate the stakeholder (or effective stakeholder surrogate) into the design process, we have the opportunity to encourage a satisfactory degree of convergence between design requirements and stakeholder expectations and demands. The terms *relevant*, *reasonable*, and *human expectations* play an extremely important role in the design of services involving human beings. We have identified these highly subjective terms by casting them into four warranty classes - express, implied, tacit, pending - based on their degree of visibility to the designer and the stakeholder.

The benefits of this work for requirements engineering are a better understanding of the notion of service and service warranties. Indeed, even direct users of IT are generally unaware of the services provided to them, the warranties they benefit from and their cost. The warranties are then taken as a given by stakeholders but are extremely expensive to maintain for IT departments. A method, which can help define service warranties that are neither too broad nor too narrow for the needs of the service stakeholder, is of potentially high value for organizations and their IT departments. Analyzing stakeholders' norms and tolerances can help define these requisite service warranties.

7. References

[1] Cornell University Law School, *U.S. Code collection*, http://www.law.cornell.edu/uscode/15/usc_sec_15_00002301----000-.html, accessed 2/10/09.

[2] GAUSE D.C., "Why Context Matters – And What Can We do about It?" *IEEE Software*, Vol. 22, No. 5, Sept/Oct 2005.

[3] Gordijn, J. and Akkermans, J.M., "Value-based requirements engineering: exploring innovative e-commerce ideas", *Requirement Engineering*, Vol. 8, No. 2, p. 114 – 134, July 2003

[4] Grandry, E., Dubois, E., Picard, M. and Rifault A., Managing the Alignment between Business and Software Services Requirements from a Capability Model Perspective, in Mähönen, P., Pohl, K., and Priol, T., (Eds.): ServiceWave 2008, LNCS 5377, pp. 171–182, Springer-Verlag 2008.

[5] Lauesen, S., *Software Requirements: Styles and Techniques*, Addison-Wesley, 2002.

[6] Office of Government Commerce, ITIL *Service Strategy*, TSO, London, 2007.

[7] Regev, G. and Wegmann, A., "Where do Goals Come from: the Underlying Principles of Goal-Oriented Requirements Engineering", *Proceedings of the 13th IEEE International Requirements Engineering Conference (RE'05)*, Paris, France, 29 August - 2 September 2005.

[8] Regev, G., Hayard, O., Gause, D.C. and Wegmann, A., "Toward a Service Management Quality Model", *Submitted 15th International Working Conference on Requirements Engineering: Foundation for Software Quality (REFSQ'09)*, Amsterdam, The Netherlands, June 8-9, 2009. In review.

[9] von Bertalanffy, L. *General System Theory*, New York: George Braziller, 1968.

[10] Wegmann, A., "On the Systemic Enterprise Architecture Methodology (SEAM)," in *Proceedings of International Conference on Enterprise Information Systems 2003* (ICEIS 2003), Angers, France.

[11] Wegmann, A., Regev, G., Garret, G-A. and Marechal, F., "Specifying Services for ITIL Service Management", in Proc. International Workshop on Service-Oriented Computing Consequences for Engineering Requirements (SOCCER'08), Barcelona, Spain, September 08, 2008.

[12] Weinberg, G. M., *An Introduction to General Systems Thinking*, Wiley & Sons. New York, 1975.

[13] Weinberg, G. M. and Weinberg, D., *General Principles of Systems Design*, Dorset House. New York, 1988.

[14] Yu, E.S.K. "Towards modelling and reasoning support for early-phase requirements engineering," in *Proceedings of the Third IEEE International Symposium on Requirements Engineering*, Annapolis, Maryland, January 1997

[15] Zeithaml, V.A., Parasuraman, A. and Berry, L.L., *Delivering Quality Service: Balancing Customer Perceptions and Expectations*, Free Press, NY, NY, 1990.