

Business-Driven Decision Support for Change Management: Planning and Scheduling of Changes

Jacques Sauvé¹, Rodrigo Rebouças¹, Antão Moura¹, Claudio Bartolini²,
Abdel Boulmakoul³, and David Trastour³

¹Departamento de Sistemas e Computação - University of Campina Grande (UFCG), Brazil

²HP Laboratories Palo Alto, USA

³HP Laboratories Bristol, UK

{jacques, rodrigor, antao}@dsc.ufcg.edu.br,

{claudio.bartolini, abdel.boulmakoul, david.trastour}@hp.com

Abstract. From the results of a web survey we carried out in 2006, the main challenge in IT change management from a change manager's perspective was identified as *planning and scheduling of changes*. This paper begins to address this problem by taking business considerations into account; this is done through a business-driven IT management (BDIM) approach. A reference architecture that follows BDIM principles is sketched; it includes a mathematical model linking IT availability metrics to business objectives. Monetary loss due to service level violations on service availability is used as the main business metric. We present a numerical illustration of how the derived metrics may support change management decisions in order to plan and schedule changes to minimize adverse business impact.

Keywords: Change management, business-driven IT management, service level management, Information Technology Infrastructure Library (ITIL), business metrics, modeling, performance evaluation, business impact, decision and negotiation support tools.

1 Introduction

IT management has become more user-centric and less service provider-dependent with the popularity of the practices recommended by the Information Technology Infrastructure Library – ITIL [3], which is used as the basis for the IT Service Management framework – ITSM [8]. ITSM defines a number of processes that are organized into 5 modules: security management; IT & communication infrastructure management; application management; service support (incident, problem, configuration, change and release management processes); and service delivery (service level, capacity, availability, continuity and financial management for IT services). Within the realm of ITSM, this paper focuses on the *change management* process.

ITSM expresses goals and gives guidelines to IT managers for ensuring smooth running of IT service delivery and support. For instance, the mission of the change

management process is defined as “[ensuring] that **standardized methods and procedures** are used for efficient and **prompt handling of all changes**, in order to **minimize the impact** of any related incidents upon service” [3]. However, it falls short of defining control objectives for IT. This shortcoming is addressed by the COBIT framework (Control Objectives for Information and related Technologies) [2]. In order to gauge the maturity and quality of IT service delivery and support activities, COBIT introduces a number of key performance indicators (KPIs) that drive the process goals, which in turn are measured by process key goal indicators (KGIs). Examples of key performance indicators for the change management process are the number of emergency changes, or the number of changes that were rolled back, in a change management context. For activities in the service delivery scope, such as service level management, metrics such as service availability and reliability are taken into account.

The first wave of management software (from the early 90’s), concentrated on monitoring availability, resource consumptions levels, etc. In the last three or four years, software tools have appeared that help with other IT management activities, in particular with help desk and IT service support. These tools provide valuable help to IT managers in making informed decisions on the actions to take to ensure the smooth running of IT processes.

However, just because the IT systems are running smoothly, it does not follow that the business that IT supports is best served by it. In order to ensure business-IT alignment, metrics should be taken into account which are more business-oriented in nature, such as cost, revenue or financial loss. This consideration is the basis for the Business-Driven IT Management (BDIM) discipline [6]. BDIM steers ITSM towards business alignment, i.e., to contribute to business results. This paper uses a BDIM approach to address change management challenges.

BDIM attempts to gauge the impact that IT has on the business and aims at rethinking IT management from this perspective. BDIM involves a new culture, tools and decision-making processes that aim to *help the business*. A complete ITSM shift to BDIM requires IT personnel or automated tools to use business metrics to gauge the QoS offered to a business user. Although BDIM has been attracting mounting research efforts, attempts at investigating the feasibility and options of spreading BDIM applications to cover ITIL management processes are still scarce. Some recent applications include incident prioritization [1], capacity planning [5], and automatic change management process [4]. Embedding results of such efforts in tools for automating decision and negotiation support is at its very beginning. This is particularly true for the case of human-assisted change management processes. This paper proposes a BDIM-based solution which could be embedded in a tool to support decision and negotiation activities in a more generic, ITSM-based change management process.

The remainder of the paper is organized into sections 2 through 7. Section 2 discusses current change management challenges as elicited by a survey performed in early 2006. Section 3 begins to address some of these challenges by describing a layered reference architecture for business-driven IT change management (BDIM-CM) solutions. Section 4 details how metrics for the BDIM-CM solution may be derived. Section 5 presents a numerical illustration on how the derived metrics may

support change management decisions in order to minimize adverse business impact. Section 6 briefly examines competitive and related work, including the few tools available on the market. In section 7 we draw our conclusions and give a preview of our further work in this space.

2 Major Challenges in Change Management

The change management process comprises four groups of activities:

- Request For Change (RFC) acceptance, classification and processing;
- approval, planning of changes;
- execution, tests and reversal of changes;
- change evaluation.

Current state-of-the-practice solutions for change management suffer from several acute problems, including the volume of changes, change complexity and inappropriate tools. It appears that the most demanding challenges faced by technical personnel in charge of change management lie in activities from the first two groups. ITIL recommends that change classification be done according to change priority and change category (which components are affected). *Priority* is set according to the business importance of an RFC relative to other RFCs; *category* is determined based on the availability of resources, risk to services and on the impact of the changes. *Planning* items include scheduling, allocation of resources, budgeting, sequencing of activities, back out plans and communication. According to results from a Web questionnaire posted in the first quarter of 2006 [7], particular attention should be devoted to planning and scheduling issues. The questionnaire respondents were ITSM practitioners, all engaged in change management (some with over 10 years of experience), from 21 companies worldwide. Seventeen of these companies already have change management processes in place (11 use ITIL and 6 adopt other practices); four are just starting to implement change management. Nine of the companies are in the business of providing IT services (including consulting), 4 are telecoms, 4 are in financial services and the others are either in government, health care or manufacturing. Seven companies have yearly revenues over US\$ 1 billion while eight make under US\$ 10 million annually. Questionnaire respondents who follow ITIL change management process recommendations ranked the first 3 most important change management challenges as being (Figure 1):

1. scheduling/planning changes (with 47 points out of a maximum 55, or over 85%);
2. high number of emergency changes (43 points or 78%); and
3. RFC scope ill-definition (40 points or 72%).

The survey also indicates that:

- adopters of an ad-hoc change management process (as opposed to a formal process like ITIL's) rank planning/scheduling as the most important challenge (80%), together with "unauthorized changes"; unauthorized changes are a problem to be expected in an ad-hoc process




Fig. 1. Most frequent problems in Change Management

- “high number of emergency changes” is ranked second (73%) together with “notification of people affected”;
- “RFC scope mal-definition” is ranked third (67%), but together with “inconsistent Configuration Management DataBase (CMDB)” (again, a possibly symptom of an ad-hoc process).

This paper contributes to addressing the most critical issues of change planning/scheduling to minimize negative impact to a service provider’s business.

Properly addressing change planning and scheduling challenges is no trivial endeavor. As commented by one of the survey respondents, “scheduling is non-trivial due to people and process problems”. Changes take place in a very dynamic environment: people become unavailable, business conditions vary and “urgent changes” may materialize. Hence, change plans and schedules have to be adjusted correspondingly. A change manager may have to build and consider several plans/schedules before a given plan is actually implemented. In an outsourcing environment, negotiating change windows with business clients is another complicating, human-dependent factor. Typically, outsourcing agreements do not provide explicit information on feasible time windows for scheduling changes that affect the associated service. Since no contractual binding exists, windows may be (and are) re-negotiated, causing re-planning and re-scheduling. The sheer volume of requests for change – RFCs – makes the scheduling exercise very complicated. As an example, the HP Managed Services organization handles 300 to 400 RFCs per weekend for a single customer. Therefore change classification and planning are currently driven by technical issues with little consideration for business needs or priorities. The solution for (re-) scheduling and elaborating such diverse plans that we begin to sketch in the next section can ease the lives of those responsible for the change management process.

3 Business-Driven Planning and Scheduling of Changes

Figure 2 depicts a reference architecture for BDIM solutions for change management, built upon a three-layer hierarchical model.

The bottom part of the figure shows objects from the IT and business environments; on the IT side, this can include, for example, RFCs and the Configuration Management Data Base (CMDB). Since a Service Level Agreement

(SLA) captures the business requirements imposed on the IT function, it sits at the boundary between the two environments. From its inputs, the lower layer uses *business-IT linkage models* to produce *IT-business linkage metrics* in the sense defined in [6]: metrics that numerically capture relationships between IT causes and their effects on business results. Examples of such metrics are: risk of adversely affecting business operations if a change is not successful and the impact of an unsuccessful or delayed change (such as potential financial loss). As will be seen in the next section’s examples, these metrics are calculated from the probability of violating an SLA, which in turn may depend on the change schedule. The solution discussed here adopts *impact* as the linkage metric and uses the probability of SLA violation to estimate it.

Linkage metrics are then fed to the middle layer – labeled *decision support* - where decisions are made and used to steer activities of the change management process (to the left of Figure 2) and/or to help negotiate change management process details – such as change windows – with the IT client (at the top of the figure – *negotiation support*).




Fig. 2. A hierarchical model for business-driven change management solutions

We propose to use the above reference architecture in a change management setting according to the following BDIM approach. By knowing details of a given RFC such as the affected components (called *configuration items* or CIs by ITIL), the past history of service levels and service level objectives (SLOs), one can calculate the probability of violating a given SLA if the change is implemented at a given time. The business-IT linkage metrics produced by the linkage layer allow one to determine the expected monetary loss that will potentially result from a given change schedule. As a result, the change manager now has numerical business impact estimates from which to choose the changes that should be implemented during a particular change window. The past comments assume a known change window time. When this time

must be negotiated with the client, impact metrics will be helpful in choosing appropriate time windows to perform changes. The business perspective introduced through the business-IT linkage model eases negotiations because arguments are presented to the client in familiar business terms (in this case, as potential financial loss).

4 Business-IT Linkage Model

This section develops a model to capture the impact of changes on the business. Consider a scenario where a service organization provides support for IT services subject to SLAs including an SLO that states a minimum availability. The client organizations use the services to process “revenue-generating sessions”; an example of such a service could be an e-commerce site where site visitors generate buying transactions during sessions accessing the service.. The provider earns a fixed fee for each successfully completed session and pays a penalty whenever the SLA is violated. Whenever a service is down but the associated SLA is not (yet) violated, the provider stops collecting fees on that service since no session can be serviced. If the SLA is violated however, besides losing the contracted fees, the provider must pay a penalty to the client at the end of the SLA evaluation period.

The business objective that we consider is to minimize the financial business loss incurred by the provider due to imperfections of supporting IT services. Changes to the IT infrastructure are “imperfections” in the sense that they may force a service to be brought down to perform the changes; as a result, changes can cause business loss. In our example, the episodes that can have an impact on the business loss are *violations of the minimum availability SLO* and *system downtime* due to changes.

- **Violation of SLO on minimum availability.** Due to penalties included as SLA clauses and, more importantly, in order not to tarnish the service provider’s image, SLA violation is frequently cited by change managers as a prime driver for decision making during change planning and scheduling.
- **System downtime.** In our example, whenever the service is down and the availability SLO is not violated, the provider stops collecting fees since no session can be serviced. This has a direct impact on the business loss.

Next, we will estimate the likelihood and extent of impacting episodes due to requested changes (subsection 4.1), and derive their impact on the business loss (4.2).

4.1 Probability of SLA Violation and Extent of System Downtime

In order to calculate business loss – in the next section – expressions for the probability of SLA violation and the extent of system downtime must be obtained.

Before formalizing the analysis, let us informally explain what we seek. Imagine that the change manager has a set of changes that may be implemented in the current SLA evaluation period and that the SLA contains an SLO on service availability. Of this evaluation period (which has duration T), duration t has already elapsed and the

change manager knows how service availability is standing up so far; in other words, the past is known. The future is *not* known but may be estimated: certain changes may be performed (or not) and they may bring down service, thus affecting the availability metric. Given the knowledge of all that has occurred in the past, the set of changes that may be considered and estimates of future availability, which changes should the change manager choose to perform in the current SLA evaluation period so as to minimize business loss?

We now formalize the problem and provide a solution. Let us first consider a single IT service s_j from the provider’s set of services, $S=\{s_1, \dots, s_{|S|}\}$ and assume that the associated Service Level Agreement (SLA) in force for s_j has a Service Level Objective (SLO) on *availability*, A_j^{min} . Let the *mean service availability* for s_j be calculated over an evaluation period T ¹, as determined by the associated SLA. This mean availability takes on a different value over each evaluation period and it is thus a random variable, denoted by \tilde{A}_j . We indicate the cumulative distribution of the service availability random variable with $F_j(x) = \Pr[\tilde{A}_j \leq x]$. Without loss of generality, let the current evaluation for availability start at time 0 and end at time T . Let us examine the situation at a point in time, t , s.t. $0 \leq t \leq T$ when the change manager must make scheduling decisions. Let the availability over period (t_1, t_2) be $A_j(t_1, t_2)$. The past mean availability over the time period $[0, t]$ is known (it is measured) and is simply: $A_j(0, t)$. The future mean availability over time period $[t, T]$ is $A_j(t, T)$. Finally the overall availability over the whole SLA evaluation period is $\tilde{A}_j = A_j(0, T)$. Now, we ask: “At time t , what is the probability that the availability threshold, A_j^{min} , specified in the SLA will be violated by time T ?” The mean availability, $A_j(0, T)$, over the whole evaluation period, $[0, T]$, can be calculated from the mean values of past and future availability by summing up the uptime over both time periods:

$$\tilde{A}_j = A_j(0, T) = \frac{uptime}{T} = \frac{A_j(0, t)t + A_j(t, T)(T - t)}{T} \tag{1}$$

In the above, the term $A_j(0, t)t$ is the uptime accumulated in the past and $A_j(t, T)(T-t)$ is the expected future uptime. Then the distribution for availability, given that time has reached t , follows:

$$\Pr[A_j(0, T) \leq x | t] = \Pr\left[A_j(t, T) \leq \frac{xT - A_j(0, t)t}{T - t} \right] \tag{2}$$

Given that the probability distribution for availability is assumed to be the same over any time period in the interval $[0, T]$, we can now express $V_j(t, T, A_j^{min})$, the probability, at time t , of violating the availability SLO for service s_j by time T :

$$V_j(t, T, A_j^{min}) = \Pr[A_j(0, T) \leq A_j^{min} | t] = F_j\left(\frac{A_j^{min}T - A_j(0, t)t}{T - t}\right) \tag{3}$$

¹For simplicity we assume that all services have this same evaluation period.

This result does not take into account the fact that, in the future, changes will be implemented, that these changes may bring down the service and that the situation for availability is actually worse than that given above. We now turn our attention to the impact of changes affecting service s_j . Let $C = \{c_1, \dots, c_{|C|}\}$ be the set of all changes to be considered by the change manager and let $T_n^c \subseteq [0, T]$ be the time interval during which a given change c_n in C is performed. Notice that implementing c_n may or may not affect s_j availability. Let service s_j be provisioned with a set, I_j^s , of components (Configuration Items or CIs, in ITIL parlance). If we let $I = \{i_1, \dots, i_{|I|}\}$ be the set of all CIs in the CMDB, then $I_j^s \subseteq I$. Each change is subject to a plan that specifies the time at which the change implementation will start, and a subset $I_n^c \subseteq I$ of the CIs that will be affected by change c_n . The plan specifies which CIs will be brought down and when, so that one can calculate the time at which service will be brought down (if it is not already down) and the time at which it will be available again. The set of all intervals during which service s_j becomes unavailable within T is given by the union of all s_j -affecting change intervals, T_n^c , i.e.,

$$T_j^s = \bigcup_{\forall n I_j^s \cap I_n^c \neq \emptyset} T_n^c \tag{4}$$

Now let ΔT_j^s (a scalar) correspond to the total time period during which s_j is unavailable, i.e., the sum (disconsidering overlaps) of the durations of all change intervals in T_j^s . Notice that all changes affecting service s_j will be implemented *after* the present moment (time t) so that any service downtime will need to be added to the ‘‘future’’ part of the evaluation period. Observing that the time period between $[t, T]$ but outside the service downtime called for in the change plan still obeys the same distribution of availability, we conclude that, with the changes planned for the current evaluation period, the probability of violating the SLA for service s_j is:

$$V_j(t, T, A_j^{\min}) = F_j \left(\frac{A_j^{\min} T - A_j(0, t)t}{T - t - \Delta T_j^s} \right) \tag{5}$$

In order to conclude the development, we need the cumulative probability distribution function, $F_j(x) = \Pr[\tilde{A}_j \leq x]$. A result from reliability theory [9] states that, when the uptime (time-to-failure) and downtime (repair times) are exponentially distributed, availability follows the two-parameter Beta distribution with parameters α and β . The mean value for availability is simply $E[\tilde{A}_j] = \alpha/(\alpha + \beta)$. α and β are chosen to match historical availability distribution data. Typical values are $\alpha = 7$ and $\beta = 0.03$, yielding 99.57% availability averaged over several evaluation periods.

4.2 Impact of SLA Violations on Business Loss

We can now turn our attention to estimating business loss due to a single service s_j . Table 1 summarizes the parameters of the provider’s revenue model.

Table 1. Provider's parameters for service s_j

Principal impact function variables	
$V_j(t, T, A_j^{\min})$	Probability of SLA violation, given the knowledge available at time t
ΔT_j^s	Service s_j unavailability period due to implementation of change selected by the change manager
Other impact function parameters	
π_j	Penalty (\$) for service s_j SLA violation
T	SLA evaluation period for s_j
γ_j	Session throughput for service s_j
σ_j	Fixed fee (\$) per successful session, for service s_j

At time t , the expected provider financial loss for service s_j during the current SLA evaluation period ending at time T is given by:

$$L_j^p(t, T) = V_j(t, T, A_j^{\min}) \cdot \pi_j + \gamma_j \cdot \Delta T_j^s \cdot \sigma_j \quad (6)$$

Notice that loss is a function of the time at which decisions are made (t) since, as time passes, the duration of the “past” (period $[0, t]$) becomes larger, and availability becomes more and more defined by what happened in the past.

In the case of multiple services $S = \{s_1, s_2, \dots, s_{|S|}\}$ supported by the provider and affected by multiple changes, the total expected provider loss at time t is:

$$Loss_s(t, T) = \sum_{j=1}^{|S|} L_j^p(t, T) \quad (7)$$

5 Numerical Illustration of BDIM Support in Change Management

Consider a scenario where a service provider offers services $S = \{s_1, s_2, s_3\}$: s_1 is a web auction service; s_2 an e-commerce service; and s_3 a database service. At time $t = \text{day } 10$, three changes appear on the change manager's desk: service s_1 is to be brought down due to two changes to two of its supporting CIs: an operating system (CI_1), whose change (c_1) is to be implemented in 3 hours; and, a DataBase Management System (CI_2), whose version upgrade (change c_2) is expected to last 2 hours. Furthermore, services s_2 and s_3 share a firewall (CI_3) and do not use CI_1 and CI_2 . A firewall change (c_3) is expected to last 4 hours. Table 2 lists SLA parameters for these three services as well as the availability situation at day 10. Let $T = 30$ days (monthly SLA evaluation).

Due to staff limitations, the change manager cannot do both $\{c_1, c_2\}$ and $\{c_3\}$ simultaneously. He must choose which set of changes to implement during the current SLA evaluation period: should changes $\{c_1, c_2\}$ or $\{c_3\}$ be done? Assume that there is no overlap possible between c_1 and c_2 . Table 2 indicates that s_1 is the service with the

greatest revenue stream (\$24/s) and has good (past) availability. The combined duration of the changes for this service (3 + 2 = 5 hours) will not cause SLA violation over the current SLA evaluation period (at $t=30$ days). On the other hand, choosing to implement c_3 will cause the service s_2 SLA to be violated, making the provider pay a \$10,000 penalty; SLA for service s_3 will not be violated. If one disregards all other change schedule influencing factors – such as political pressure from clients, concerns with provider image, change roll-back problems and the cost of not executing a given change with provider image, change roll-back problems and the cost of not executing a given change with provider image, change roll-back problems and the cost of not executing a given change with provider image – and if one simply analyzes SLA clauses and short-term past history, one may be tempted to opt for implementing $\{c_1, c_2\}$, affecting s_1 since this option is likely to yield a smaller loss. This is an approach commonly employed by change managers (avoid SLA violations!).

Table 2. Service configuration

Input	Service s_1	Service s_2	Service s_3
γ_j (session/s)	16	12	13
σ_j	\$ 1.5	\$ 0.7	\$ 0.8
π_j	\$ 30,000	\$ 15,000	\$ 10,000
A_j^{\min}	0.99	0.99	0.99
$A_j(0, \text{day } 10)$	0.998	0.992	0.9998

Let us now examine the situation in the light of business loss metrics (Figure 3). This figure compares business loss for both alternatives (choosing $\{c_1, c_2\}$ or $\{c_3\}$) as the decision time point changes (horizontal axis). This figure shows that, at day 10, selecting change $\{c_3\}$ is preferable since it causes lower business loss (around \$14,000 compared to \$21,000). However, should the three changes land on the change manager’s desk for decision on day 17, say, it would be more advisable to select $\{c_1, c_2\}$ since expected loss for change $\{c_3\}$ has now increased substantially (to \$27,000). The reason for this drastic change around day 16 is that, from that time on, change $\{c_3\}$ will cause an SLA violation for service s_3 . Change management decisions are dynamic in nature and our business-IT linkage metrics captures this dynamic behavior.




Fig. 3. Expected losses with changes in illustration scenario

6 Related Work

IT management software tools available on the market (such as HP OpenView ServiceDesk and ServiceCenter and BMC Routes-to-Value Change and Configuration Management) provide administrative support to the change management process by tracking a change in all phases of its lifecycle, coordinating its different activities, assigning activities to the appropriate people and monitoring its progress until it is closed. However, these tools provide no support to the decision-making process, and, although concepts such as *risk* and *impact* are present, their definition is rather ambiguous and their assessment is left to the tool user.

CHAMPS [4], a research prototype out of IBM Research, represents the state of the art in automation for change management, but it does not address aspects of project management of the changes such as scheduling activities that require human intervention. Further, it is assumed that business impact is an input parameter and the intended application is to an autonomic computing setting. The linkage model in our work helps to evaluate business impact. Our solution addresses change management challenges holistically: it considers all three components people, process and technology. Thus, our work may be seen as complementary to that of [4]. The work in [1] brings ideas that could be adapted for prioritization/classification of RFCs, since RFCs are frequently related to incidents. Usage of utility functions is particularly attractive.

7 Conclusions and Future Work

From the results of a web survey we carried out in earlier 2006 [7], the main challenges in IT change management were identified as 1) planning/scheduling changes, 2) high number of emergency changes and 3) ill-definition or wrong scoping of requests for change. In this paper we have begun to address the problem of planning and scheduling changes by taking business considerations into account, following a business-driven IT management (BDIM, [6]) approach. We have sketched a reference architecture that follows BDIM principles; the architecture includes a model linking IT availability metrics to business objectives (in our example: minimizing financial loss due service unavailability and SLA violations). A numerical illustration was presented to show how the derived metrics may support change management decisions in order to plan and schedule changes to minimize adverse business impact.

This initial work supports the conception of an automated tool for decision support for planning and scheduling changes. We have received encouraging feedback from the respondents of our survey in [7], to whom we presented the scenario here exposed in a follow up interaction. The respondents agree that the information these metrics provide will definitely “add value to the decision process”. However, much remains to be done before we can embody the capabilities here described into a software tool that is complete enough to be of value to change managers in their decision making activities and while negotiating with customers. The main contribution of this paper is the formalization of a sound base for supporting change scheduling and planning. Our next step will be to formally bring in change windows and solve an optimization

problem to find the “best” allocation of changes to change windows. Design and implementation of the decision support tool is also the subject of a future phase of our research.

Acknowledgements

This work was developed in collaboration with HP Brazil R&D.

References

1. Bartolini, C., and Sallé, M., “Business Driven Prioritization of Service Incidents”, *In Proceedings of the 15th IFIP/IEEE International Workshop on Distributed Systems: Operations and Management (DSOM 2004)*, 15-17 Nov. 2004, Davis, CA, USA, pp 64-75.
2. IT Governance Institute, “Cobit 4th Edition”, 2006, www.isaca.org/cobit.htm
3. IT Infrastructure Library, “ITIL Service Delivery” and “ITIL Service Support”, Office of Government Commerce, UK, 2003.
4. Keller, A., Hellerstein, J., Wolf, J.L., Wu, K. and Krishnan, V., “The CHAMPS System: Change Management with Planning and Scheduling”, *In Proceedings of the IEEE/IFIP Network Operations and Management Symposium (NOMS 2004)*, IEEE Press, April 2004, pp. 395-408.
5. Sauvé, J. P., Marques, F. T., Moura, J. A. B., Sampaio, M. C., Jornada, J., Radziuk, E., Optimal Design of E-Commerce Site Infrastructure from a Business Perspective, In: *Proceedings of the 39th Annual Hawaii International Conference on System Sciences (HICSS'06)*, Waikoloa, Hawaii. IEEE Computer Society, 2006. v.8. p.178.3 - 178.3
6. Sauvé, J., Moura, A., Sampaio, M., Jornada, J. and Radziuk, E., “An Introductory Overview and Survey of Business-Driven IT Management”, in *Proceedings of the 1st IEEE / IFIP International Workshop On Business-Driven IT Management*, in conjunction with NOMS 2006, Vancouver, Canada, pp. 1-10.
7. The HP-Bottom Line Project, “IT Change Management Challenges – Results of 2006 Web Survey” Technical Report DSC005-06, Computing Systems Department, Federal University of Campina Grande, Brazil, March 2006. http://www.bottomlineproject.com/bl/_media/techreport/005-2006.pdf
8. Van Bon, J., Chief Editor, “IT Service Management, an introduction based on ITIL”, itSMF Library, Van Haren Publishing, 2004.
9. Abramowitz, M. and Stegun, I. A. (Eds.). "Beta Function" and "Incomplete Beta Function." §6.2 and 6.6 in *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables, 9th printing*. New York: Dover, pp. 258 and 263, 1972.