# Business-impact analysis and simulation of critical incidents in IT service management

C. Bartolini
HP Labs
1501 Page Mill Rd
Palo Alto, CA 94304 USA

C. Stefanelli, M. Tortonesi
University of Ferrara
Via Saragat, 1
44100 Ferrara, Italy

*Abstract*—**Service disruptions can have a considerable impact on business operations of IT support organizations, thus calling for the implementation of efficient incident management and service restoration processes. The evaluation and improvement of incident management strategies currently in place, in order to minimize the business-impact of major service disruptions, is a very arduous task which goes beyond the optimization with respect to IT-level metrics. This paper presents HANNIBAL, a decision support tool for the business impact analysis and improvement of the incident management process. HANNIBAL evaluates possible strategies for an IT support organization to deal with major service disruptions. HANNIBAL then selects the strategy with the best alignment to the business objectives. Experimental results collected from the HANNIBAL application to a realistic case study show that business impact-driven optimization outperforms traditional performance-driven optimization.**

*Index Terms*—**Business-driven IT management (BDIM), decision support, Information Technology Infrastructure Library (ITIL), IT service management, incident management.**

## I. Introduction

INCIDENT Management, as defined by the IT infrastructure library (ITIL, [1]), is the process through which IT support organizations manage to restore normal service operation after a disruption.

IT support organizations are usually very large and complex, composed of many support groups, in turn organized into several support levels. Their complexity hinders both the evaluation of currently adopted incident management strategies and the improvement of internal processes in order to provide the most effective response to service disruptions.

IT support organizations typically evaluate their performance through IT-level metrics, such as latency and throughput in the incident response process. Organizations then use these metrics as benchmark values to improve their performance in service recovery operations. Such performance metrics, however, fail to capture the whole impact of service disruption on IT support organizations, as they do not consider the effects on business. For this reason, business-level metrics or *Key Performance Indicators (KPIs)* are used, based in turn on Service Level Objective (SLO) violations and penalties.

KPIs allow the IT support organization to assess more precisely the impact of service disruption on business. In fact, while some incidents may have limited or no consequences on business critical services, others, called *critical incidents*, can induce SLO violations, thus causing major service disruptions, and can have a considerable impact on business operations. The full extent of the effect of critical incidents on IT support organizations cannot be captured through IT-level metrics alone.

For effective optimization of the incident management process, it is indispensable to prioritize the IT intervention to minimize the business impact of critical incidents. To this end, *business impact analysis* represents the best criterion to follow for the improvement of the incident management process. Business impact-driven optimization aims at minimizing the adverse impact of service disruptions on the business, by considering all the costs attached to critical incident occurrences.

The complexity of both the IT service management domain and the incident management process, makes it very hard to tackle the business impact-driven optimization problem analytically, and calls for automated decision supports. In this context, this paper presents HANNIBAL, a decision support tool for improving the organizational, structural, and behavioral components of IT support organizations in dealing with critical incidents, according to business impact analysis considerations. More specifically, HANNIBAL analyzes a set of alternative strategies for managing critical incidents and it helps a user to select the one with the best alignment with business objectives.

HANNIBAL has been tested in a case study representative of the complexity of real-life world-class IT support organizations. The collected experimental results show that business impact-driven optimization is significantly different from performance-driven optimization, and allows a better alignment with business objectives.

The remainder of the paper is structured as follows. Section II provides a brief overview of the IT incident management process and describes the importance of strategic consideration for improving the effectiveness of IT support organizations. Section III introduces the HANNIBAL decision support tool and Section IV describes its architecture and implementation. Section V presents the experimental results. Section VI reviews the related works and Section VII provides conclusive remarks and future work considerations.

## II. Incident Management

The IT Infrastructure Library (ITIL, [1]) divides the incident management process into several steps: incident detection and recording, classification and initial support, investigation and diagnosis, resolution and recovery, closure, and tracking.

A typical IT support organization is structured as a network of support groups, each comprising a set of technicians. Technicians, and (as a consequence) the support groups they are assigned to, have a definite work schedule, often depending on geographical residence. Support groups are divided into *support levels* (usually three to five), with lower level groups dealing with generic issues and higher level groups handling technical and time-consuming tasks. Support groups are further specialized by category of incidents that they deal with (network, server, etc…).

The *Help Desk* function of a support organization represents the interface for customers reporting IT service disruptions. In response to a customer request, the Help Desk "opens" an incident, sometimes called *trouble-ticket* or simply ticket. The incident is then "assigned" to a specific support group. The technicians in the support group work on restoring disrupted services and may be able to "close" the incident. More often, they will see the need to "reassign" it to a different support group (usually escalating to a higher support level). As a result, an incident will have different states and will be handled by different support groups throughout its lifetime. At each of these steps, the ticket is updated with the pertinent information, such as current state and related service restoration activity. If, for some reason, a customer requests the organization to stop working on an incident, the incident is placed in a "suspended" state to avoid incurring into SLO *(Service Level Objective)* penalties. Once the disruption is repaired, the ticket is placed in "closed" state until the end-user confirms that the service has been fully restored. In this case, the incident is "resolved" and its lifecycle ends.

At their arrival, incidents are assessed for importance, severity and urgency, and assigned a *priority level* for resolution. Because the incidents represent service disruptions that have an adverse impact on the business, effective incident prioritization has substantial bearing on the effectiveness of the IT support organization in dealing with incidents.

Among other things, incident priority levels are used for determining the order in which support groups deal with the incidents. Because support groups are dimensioned to deal with the expected traffic rate of incidents, they keep a queue of incidents yet to be dealt with. As technicians become available to work on incidents, they retrieve incidents from the queue according to a pre-determined *policy*. During incident lifetime, incident priority levels may be changed several times.

Since IT support organizations are subject to frequent changes with regards to service implementation and restoration practices, incident management strategies must be periodically evaluated and changed when deemed obsolete. However, the complexity of IT support organizations is an obstacle to effective verification of the alignment of current organizational, structural, and behavioral processes with the strategic objectives defined at the business management level.

In fact, the performance assessment of the incident management function is a very complex procedure which involves the impact evaluation of the current incident management strategy on a set of metrics, such as Mean Incidents Closed Daily (MICD) and Mean Time To (incident) Resolution (MTTR) [2, 3].

Performance analysis and optimization are also organization-specific procedures, since the impact of service disruptions, and consequently the metrics to consider, vary with the nature of the services and the types of disruptions that occur.

In addition, incident management processes can have a non-negligible impacts on business, e.g., through SLO violation penalties. As a result, there is the need to consider business impact analysis in the optimization of the incident management process.

Of particular interest is the case of incidents that can bring to SLO violations (or *critical incidents*). These incidents are the most important as they can have a considerable impact on business operations.

## III. HANNIBAL: A Business Impact-driven Decision Support Tool for Critical Incident Management Strategies

Incidents that can bring to SLO violations, or *critical incidents*, are of particular interest for IT support organizations. In fact, while they are not as frequent as normal (non-critical) incidents, critical incidents have the most significant impact on business operations.

This suggests the implementation of special strategies for the management of critical incidents, optimized to reduce their business impact. In fact, the performance of an IT support organization in the incident management process is subject to dramatic variation depending on the effectiveness of critical incident routing and on the efficiency of each single support group in dealing with tickets.

As a result, the optimization of critical incident management should consider both strategies increasing the *effectiveness* of incident routing (comparing alternate support group policies for forwarding and escalating incidents) and strategies increasing the *efficiency* in dealing with the most important incidents first (comparing alternate prioritization policies for extracting incidents from support group queues).

However, the assessment of both the *tangibles* (immediately visible costs due to SLO penalties, hardware and personnel) and the *intangibles* ("hidden" costs due to increase/reduction of incident response efficiency) business impact factors of various strategies is a very complex process, and calls for support tools to enable informed and accurate decision making.

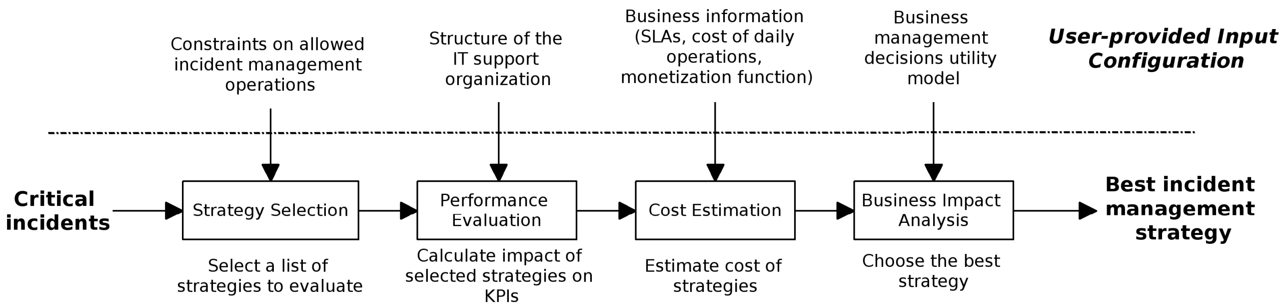


Fig. 1: The Hannibal Decision Support Process.

HANNIBAL is a business impact-driven decision support tool for the selection of strategies in critical incident management. HANNIBAL enables business managers to make well informed decisions about the critical incident management processes, at the organizational, structural, and behavioral level.

The HANNIBAL decision support tool embodies a 4-phase process, as described in Figure 1. Starting from a set of critical incidents, HANNIBAL proposes a set of candidate strategies (*Strategy Selection* phase), evaluates the performance of the incident management process in the context of each candidate strategy (*Performance Evaluation* phase), calculates the cost of each candidate strategy (*Cost Estimation* phase), and finally reports which strategy has the least business impact (*Business Impact Analysis* phase).

HANNIBAL also integrates reporting functions, which provide a detailed analysis of strategy evaluation process, thereby offering a comprehensive set of information to support business decisions.

The rest of this section discusses in details the 4 phases of the HANNIBAL decision support process: *Strategy Selection*, *Performance Evaluation*, *Cost Estimation*, and *Decision Making*.

### A. Strategy Selection

In the Strategy Selection phase, HANNIBAL considers several strategies for critical incidents management. The strategies taken into account cover the whole space of allowed options, according to the user-specified constraints.

Before proceeding to the next phase, a user can modify the set of selected strategies for evaluation. For instance, users might want to further refine or restrict subset of the proposed incident management strategies, and/or add new candidate strategies to be considered in the evaluation process. There are two main performance aspect to take into consideration: *support level efficiency* and *incident routing effectiveness*.

In order to address *support group efficiency*, HANNIBAL considers different policies for critical incident prioritization. The policies are represented by criteria over fundamental attributes of the incidents, in particular incident priority levels. Example of such criteria are: "*if the priority level is* low priority*, when the incident is escalated to another support group, it is put at the end of the incoming incident queue*"; "*if the priority level is* high priority*, when the incident is escalated to another support group, it is put at the front of the incoming incident queue*"; "*if the priority level is* very high priority*, when the incident is escalated to another support group it is immediately assigned an operator, thereby preempting other incidents*", and so on.

In order to address *incident routing effectiveness*, HANNIBAL also considers different incident routing strategies, such as assignment to specific support groups or to several operators at the same time. More specifically, the support group routing strategies include assignment of critical incidents to the support group with the *shortest incident response* (shortest average time spent by incidents waiting on incoming incident queue), to the support group with the *lowest load* (largest operator idle/busy time ratio), to the support group which provides the *best skills* for incident resolution (more skilled operator), to the *largest* support group (largest operator set), to the most appropriate support group taking into account *geographies and time zones*, or to a *random* support group. In the context of the selected support group, different operator assignment strategies are considered, such as assignment to the *best skilled* operator, and assignment to the *first available* (random) operator.

### B. Performance Evaluation

In the Performance Evaluation phase, HANNIBAL estimates the impact of the strategies selected at the previous phase on the IT support organization performance. In particular, this phase evaluates the impact of strategies on the KPIs.

The performance evaluation task is performed via *what-if scenario analysis.* To this end, HANNIBAL leverages on the SYMIAN tool to reenact IT support organization processes [4]. SYMIAN implements an accurate model of IT organizations behavior, enabling the evaluation of their incident management performance.

### C. Cost Estimation

In the Cost Estimation phase, HANNIBAL calculates the cost of implementing the strategies under evaluation in the IT support organization. The calculation has to consider several factors: the costs for the strategies implementation itself, the costs related to SLO violations occurred in the context of the strategies, and variations of normal operations costs due to strategy implementation.

The costs directly related to the strategies implementation depend from the specific strategies considered. For instance, the cost of adding new operators to a specific support group must consider the costs for operator training, equipment, and salary. The cost of implementing software/hardware replacement and/or upgrades, instead, must consider the cost for buying new software/hardware and the cost for installation, configuration, and training. As a result, HANNIBAL requires the user to provide specific implementation costs for each strategy to be evaluated.

HANNIBAL also considers SLO violation penalties due to strategies implementation. To this end, HANNIBAL requires users to define the conditions in which SLO violations occur and their penalty amount. The tool then uses the values of service level indicators obtained from the previous phase to find whether SLO violations occur in the context of the strategies under evaluation, and calculates their costs.

### D. Business Impact Analysis

In the Business Impact Analysis phase, HANNIBAL calculates the business impacts of strategies and ranks them according to user-provided business management preferences.

Business objectives are the targets of the incident management process, defined at the business management level. HANNIBAL models business objectives following the MBO information model [3]. A business objective is defined by defining a *target region* (usually expressed through a constraint) over a *KPI* (key performance indicator). A business objective is met when the specific KPI value lies within the target region at the end of the evaluation period. Each business objective has an associated *weight* - which is configurable by the HANNIBAL user - and expresses user preferences in terms of relative importance of business objectives. As thoroughly described in [3], the MBO information model for business objectives used by HANNIBAL models business management preferences using weights based on the Balanced Scorecard concept [5]. Importance weights are associated to perspectives of the business scorecard, and in turn these are further modified by weights associated to the objective proper. Table I provides examples of business objectives, their perspectives, and related importance weights.

HANNIBAL computes the *alignment* of all the evaluated strategies with the user-provided business objectives to the find the strategies with the minimum business impact.

TABLE I
EXAMPLE OF BUSINESS OBJECTIVES.

| Business objective | KPI | Target region | Obj. Wt. | Final Wt. |
|---|---|---|---|---|
| *External Perspective – Importance weight: 0.3* | | | | |
| *Customer Satisfaction* | Total Number of SLO violations | Less than 10 violations | 1.0 | 0.30 |
| *Financial Perspective – Importance weight: 0.7* | | | | |
| *Cost of implementing new strategies* | Total cost of implementing new strategies | lower than 50,000 $ per three month period | 0.6 | 0.42 |
| *Aggregated cost for SLO penalties* | Total Cost of SLO penalties | lower than 10,000 $ per month | 0.4 | 0.28 |

## IV. HANNIBAL ARCHITECTURE AND IMPLEMENTATION

HANNIBAL implements each of the four phases of the decision support pipeline through a specific component. These are the *Strategy Selector* component, the *SYMIAN* tool for what-if scenario analysis, the *Analyzer* component, and the *Alignment Estimator* component.

The Strategy Selector component implements the selection of strategies according to user-provided configuration and constraints. The output of this phase represents the whole space of candidate strategies to consider for the evaluation.

SYMIAN is the component realizing the performance evaluation phase. SYMIAN implements an accurate model of IT support organizations which allows, via discrete event simulation, to reproduce their behaviour and to evaluate their KPIs in the context of each candidate strategy for critical incident management. For more information on the SYMIAN implementation, see [4].

The Analyzer component implements both performance analysis and the cost analysis of strategies on the incident management process.

The Alignment Estimator component calculates the alignment of strategies with business objectives, and compares them to find out which one has the minimum impact on business.

The components implementing the decision support process are supported by other components, implementing coordination and auxiliary functions.

The *Coordinator* component directs and supervises the decision support process. It connects the decision support components together, performing ad-hoc transformation and processing of data when necessary.

The *User Interface* component allows users to define configuration parameters, to launch the decision making process, and to save its outcomes to file. The User Interface component provides both an interactive textual and a non-interactive command-line interface.

The *Configuration Manager* takes care of all the aspects regarding HANNIBAL configuration. Among the functions provided by the Configuration Manager are configuration file parsing and validation of configuration parameters.

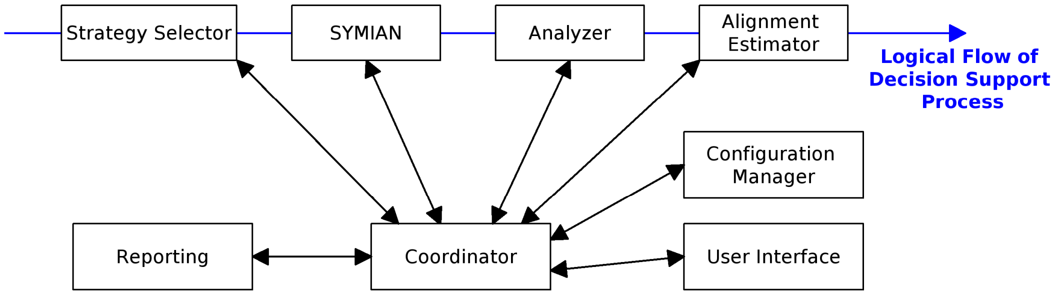

Figure 2. Architecture of the Hannibal decision support tool.

The *Reporting* component provides reporting functions in order to offer a comprehensive set of information to support business decisions. The Reporting component implements statistical analysis functions and integrates with the Gnuplot data visualization tool (http://www.gnuplot.info/) to plot time-varying values such as incoming incident queue sizes at the support groups.

HANNIBAL is implemented in the Ruby programming language (http://www.ruby-lang.org/). Ruby was chosen for its excellent support for rapid application development, by means of its remarkable extensibility and its capabilities to define domain-specific languages, and for the availability of a wide range of high-quality scientific libraries and tools.

The Ruby language is particularly well suited for the implementation of discrete event simulation-based tools like HANNIBAL, and offers satisfying performance levels which allow HANNIBAL simulations with a volume of incidents up to several tens of thousand to run in a few minutes.

## V. Experimental Results

This section presents an experimental evaluation of the HANNIBAL effectiveness in the analysis and optimization of the incident management process. Here, HANNIBAL is applied to optimize a case study IT support organization modeled according to real-life experiences. The experiments compare gains in performance when optimization is driven by business impact consideration rather than by IT level metrics.

### A. *Experiment Configuration*

The IT organization subject of this experimental evaluation, Wolfe Inc., consists of a help-desk plus 2 support levels (levels 0-2), and 31 support groups, comprising of 348 technicians.

To limit the complexity of the case study, the organization model assumes the routing of incidents in the Wolfe Inc. organization to be unidirectional, meaning that support groups of level N only receive incidents from support groups of level N-1 and escalate incidents to support groups of level N+1.

The experiments covered one month (31 days) of simulated time, starting from August 12$^{th}$ 2008, 11:10AM UTC. The first day of simulated time was used exclusively to prime the simulation environment in order to prevent "cold start" measurements from affecting the simulation accuracy, and as such was not considered for the performance evaluation.

Wolfe Inc. deals with both non-critical and critical incidents. Non-critical incidents arrive continuously at a swift pace and need only a limited amount of work time before they can be closed, while critical incident arrive rarely and require a significant amount of work time for service restoration.

In all the experiments, non-critical incidents have random arrival times following an exponential probability distribution with a mean arrival rate of 30 minutes. The exponential probability distribution models non-critical incident arrival as a *memory-less* process. Non-critical incidents require a random amount of work at every support level, modeled according to a uniform distribution probability in the (0, 30 minutes] time interval.

The 3 critical incidents considered in the experiments were instead deterministically modeled, with regards to their arrival time and time to resolution at each support level. Table II provides their detailed characterization.

The set of possible strategies considered in the optimization process was intentionally kept small, to demonstrate more clearly that the performance-driven and business impact-driven optimization processes lead to very different results. More specifically, two policies were considered for incident prioritization: *high priority* and *normal priority*, and two strategies for support group assignment: *random* and *largest*. For operator assignment, only the *first available* policy was considered.

Table III shows the four SLOs considered in these experiments. The first SLO imposes a maximum MTTR value of 10 hours. The second SLO states that no more than 530 incidents per month shall be closed in more than 2 hours from the time of their arrival. The third SLO states that all critical incidents shall be closed within one month (which coincides with the end of the simulation period). If any one of these SLOs is not met, the Wolfe Inc. organization is charged a penalty of 10,000 $ and a SLO violation is reported. The fourth SLO established an extra penalty of 10,000 $ if more than two violations occur.

The business objectives considered for the optimization are presented in Table IV. Only two objectives were taken into account. The most important objective - with a weight of 0.65 - aims at keeping the total cost of monthly SLO penalties under 10,000 $. The secondary objective - with a weight of 0.35 - is a measure of customer satisfaction, limiting the number of allowed SLO violations to two.

TABLE II
CHARACTERIZATION OF CRITICAL INCIDENTS CONSIDERED IN SIMULATIONS.

| Critical Incident | Arrival time (from simulation start time) | Required work time for resolution |
|---|---|---|
| 1 | After 12 days and 5 hours | L0: 30 minutes<br>L1: 1 day<br>L2: 2 days and 12 hours |
| 2 | After 16 days and 9 hours | L0: 45 minutes<br>L1: 1 day and 6 hours<br>L2: 2 days |
| 3 | After 20 days and 11 hours | L0: 50 minutes<br>L1: 16 hours<br>L2: 1 day and 12 hours |

TABLE III
CHARACTERIZATION OF SERVICE LEVEL OBJECTIVES.

| Condition | Penalty | Triggers violation |
|---|---|---|
| MTTR metric greater than 10 hours | 10,000 $ | Yes |
| Time to closure greater than 2 hours for more than 530 incidents per month | 10,000 $ | Yes |
| All critical incidents should be closed in the simulation period | 10,000 $ | Yes |
| Violations > 2 | 10,000 $ | N/A |

Before the optimization process, a performance evaluation of the WOLFE INC. organization in case of no critical incident occurrence was conducted. The purpose of this evaluation is to obtain a benchmark serving as a comparison for the optimization outcome.

Table V provides the values for the Mean Incidents Closed Daily (MICD) and Mean Time To (incident) Resolution (MTTR) metrics obtained from the simulation. The Mean Wait Time (MWT) metric, defined as the mean time spent in queues waiting for an available operator per closed incident, and the MWT/MTTR ratio are also provided as an indication on the efficiency of service restoration operations.

### B. IT-driven Optimization

The first experiment optimizes performance using minimization of service disruption time as objective. Here HANNIBAL is configured to ignore business objectives and select the strategy which minimizes the MTTR metric.

TABLE IV
CHARACTERIZATION OF BUSINESS OBJECTIVES.

| Business Objective | KPI | Target region | Wt |
|---|---|---|---|
| *Customer Satisfaction* | Total Number of SLO violations | At most 2 violations | .35 |
| *Aggregated cost for SLO penalties* | Total Cost of SLO penalties | At most 10,000 $ per month | .65 |

TABLE V
PERFORMANCE ANALYSIS OF THE WOLFE INC. IT SUPPORT ORGANIZATION WITHOUT CRITICAL INCIDENTS OCCURRENCE.

| | |
|---|---|
| Total incidents arrived | 661 |
| Incidents arrived after warm-up threshold | 645 |
| Closed incidents | 632 |
| MICD | 21.07 |
| MTTR | 10 hours, 21 minutes, and 41 seconds |
| MWT | 9 hours, 37 minutes, and 37 seconds |
| MWR/MTTR ratio | 0.93 |

Table VI shows the HANNIBAL outcomes for the performance-driven optimization process: the number of incidents arrived, considered, and closed, the MTTR and MWT metrics, the characterization of the selected strategy for critical incidents, the SLO penalties and violations occurred, and finally the strategy alignment with business objectives.

Analyzing the data presented in Table VI, it is possible to notice that the 3 critical incidents do not significantly impact the MTTR metric (only a 6% increase compared with the reference value in case of no critical incident occurrence), and that the selected strategy satisfies the first business objective, as only 2 SLO violations occurred. However, the total amount of SLO penalties was well above the 10,000 $ threshold set by the second business objective. As a result, the strategy selected by HANNIBAL has only a rather low (35%) value of alignment with the given business objectives.

### C. Business impact-driven Optimization

The objective of the business impact-driven optimization process is the selection of the strategy for critical incident management scoring the highest level of alignment with the given business objectives. Table VII shows the HANNIBAL outcomes for the business impact-driven optimization process.

TABLE VI
RESULTS OF PERFORMANCE-DRIVEN OPTIMIZATION.

| | |
|---|---|
| Total incidents arrived | 664 |
| Incidents arrived after warm-up threshold | 648 |
| Closed incidents | 635 |
| MICD | 21.16 |
| MTTR | 10 hours, 59 minutes, and 35 seconds |
| MWT | 9 hours, 55 minutes, and 16 seconds |
| MWT/MTTR ratio | 0.90 |
| Selected strategy for critical incidents | incident prioritization => high,<br>supp. group assignment => largest,<br>oper. assignment => first available |
| SLO penalties | 20,000 $ |
| SLO violations | 2 |
| Alignment | 35% |

TABLE VII
RESULTS OF BUSINESS IMPACT-DRIVEN OPTIMIZATION.

| | |
|---|---|
| Total incidents arrived | 664 |
| Incidents arrived after warm-up threshold | 648 |
| Closed incidents | 636 |
| MICD | 21.2 |
| MTTR | 11 hours, 22 minutes, and 28 seconds |
| MWT | 10 hours, 18 minutes, and 10 seconds |
| MWT/MTTR ratio | 0.91 |
| Adopted strategy for critical incidents | incident prioritization => high, supp. group assignment => random, oper. assignment => first available |
| SLO penalties | 10,000 $ |
| SLO violations | 1 |
| Alignment | 100% |

The analysis of the data in Table VII shows that the selected strategy has a significant impact on the MTTR metric (a 9,8% increase compared with the reference value in case of no critical incident occurrence). However, both the SLO violations and penalties are below the thresholds set by business objectives. As a result, the strategy selected by HANNIBAL has the maximum (100%) value of alignment with the given business objectives. This case is representative of the fact that optimizing for IT metrics does not necessarily result in the best possible business performance.

## VI. RELATED WORK

This paper contributes to the research domain of Business-Driven IT Management (BDIM) [6]. Research on BDIM concerns the quantitative evaluation of interdependencies between business performance and IT solutions in order to improve the quality of IT services and the business results that IT supports.

Most of the research in BDIM (applied to change management, [7, 8, 9], capacity management [10, 11, 12], network security [13], and network configuration management [14]) has been fairly limited to the *tools* and *technology* dimensions of IT management, focusing on the fine tuning of IT infrastructure configuration and on automation as means to obtain savings and improve the efficiency of IT management processes.

The present work differs from most of the mainstream efforts in BDIM in two important ways. First, HANNIBAL's suggestions are based on longer term policies and strategies, rather than with short term optimizations of IT based parameters. Accordingly, this work can be classified along with the Tyche SLA-based management [15], and Sauve et. al's SLA design from a business perspectives [12]. Second, and more importantly, it focuses on the other two fundamental dimensions of IT management: *people* and *processes*.

Only recently have research interests on the people and processes dimensions of IT management emerged, and thus there are relatively few works in the academic literature. The IT Service Management (ITSM) practices recommended by the ITIL [1] framework include guidelines and best practices to effectively manage IT and bring value to the business. IT management and governance so far have been a subject studied mostly in the business schools circles [16]. However, the efforts in the academic research community around computer science in general, and network and systems management in particular directed at achieving business-IT integration covering all ITIL management processes, and at meeting business requirements for effective IT governance, are still at the very early stage.

Diao et al.'s work [17, 18] can serve as a representative research work investigating relationships between people, processes, and technological optimization and the impact of automation and process complexity on labor cost. They have undertaken a research effort addressing the very important question of when does it make sense to automate processes based on metrics of process complexity [17, 18]. Compared with Diao et al.'s research, the present work differs for its focus on the achievement of significant performance improvements in IT support organizations through decision support and simulation techniques.

The authors' contribution to the people and processes side of the IT service management research goes back to the study of incident management strategies, introducing the *IT management by business objectives* (MBO) methodology [2, 3]. MBO moved from the definition of business-level objectives such as those commonly used in balanced scorecards. The work presented here shares with MBO a common information model to express business objectives, key performance indicators and other fundamental concept, such as a quantitative definition of alignment with business objectives. However, with respect to [2, 3], this paper follows a novel approach that proposes and implements a detailed model of the inner functioning of the IT support organization to support what-if scenario analyses. Another important difference with the MBO methodology is that in MBO, alignment with business objective was estimated through statistical methods to drive management options, whereas the approach presented here uses simulation. The result of this choice is that quantitative alignment with business objectives is predicted through collecting data from repeated simulations, resulting in much more precise estimates of future KPI values, even when compared with time series forecast models such as Box and Jennings' ARIMA that were suggested as the basis of the MBO methodology.

The analysis of the incident management process and the IT support organization model presented in this paper build on the same foundations as the authors' work presented in [19]. However, this paper proposes richer and more accurate models, reaching far beyond the definition of metrics for the IT support performance assessment that was conducted in [19].

A central component of the system described in this paper is

SYMIAN, a decision support tool enabling performance evaluation of IT support organizations through what-if scenario analysis that was introduced in [4]. However, in [4] the optimization of the incident management process was driven by IT-level metrics. As thoroughly argued in the validation section, optimizing for IT metrics does not necessarily result in the best possible business performance.

## VII. Conclusions and Future Work

This paper introduces HANNIBAL, a decision support tool for business impact analysis and improvement of the incident management process. Experimental results collected from the HANNIBAL application to a realistic case study show that business impact-driven optimization outperforms traditional performance-driven optimization.

The significant potential demonstrated by the HANNIBAL decision support tool calls for further study, including a comprehensive and more quantitatively accurate validation of the tool effectiveness when applied to real life data.

Future versions of HANNIBAL will also consider critical incident management strategies based on the restructuring of the IT support organization by merging and splitting of support groups and re-staffing by adding, removing, or reassigning technicians to support groups.

## References


[1] The IT Infrastructure Library.

[2] C. Bartolini, M. Sallé, "Business Driven Prioritization of Service Incidents", In *Proceedings of 15th IFIP/IEEE International Workshop on Distributed Systems Operations and Management (DSOM 2004).*

[3] C. Bartolini, M. Sallé, and D. Trastour ,"IT Service Management driven by Business Objectives - An Application to Incident Management", in *Proceedings of 10th IEEE/IFIP Network Operations and Management Symposium (NOMS 2006)*, 3-7 April 2006, Vancouver, Canada.

[4] C. Bartolini, C. Stefanelli, and M. Tortonesi, "SYMIAN: a Simulation Tool for the Optimization of the IT Incident Management Process", in *Proceedings of 19th IFIP/IEEE International Workshop on Distributed Systems: Operations and Management (DSOM 2008)*, 25-26 September 2008, Island of Samos, Greece.

[5] R. S. Kaplan, D. P. Norton, "The Balanced Scorecard: Translating Strategy into Action" Harvard Business School Press, 1996.

[6] A. Moura, J. Sauvé, and C. Bartolini, "Research Challenges of Business–Driven IT Management", *In Proceedings of the 2nd IEEE / IFIP International Workshop On Business-Driven IT Management (BDIM 2007)*, Munich, Germany.

[7] 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.

[8] Sauvé, J., Rebouças, R., Moura, A., Bartolini, C., Boulmakoul, A., Trastour, D., "Business-driven decision support for change management: planning and scheduling of changes", In Proceedings of the DSOM 2006, Dublin, Ireland.

[9] Trastour, D., Rahmouni, M., Bartolini, C., "Activity-Based Scheduling of IT Changes" In Proceedings of First International Conference on Autonomous Infrastructure, Management and Security, AIMS 2007, Oslo, Norway, 73-84.

[10] Aiber, S., Gilat, D., Landau, A., Razinkov, N., Sela, A. and Wasserkrug, S. "Autonomic Self–Optimization According to Business Objectives"; In Proceedings of the International Conference on Autonomic Computing, 2004.

[11] Menascé, D., Almeida, V.A.F., Fonseca, R. and Mendes, M.A., "Business-Oriented Resource Management Policies for e-Commerce Servers", Performance Evaluation 42, Elsevier Science, 2000, pp. 223-239.

[12] Sauvé, J., Marques, F., Moura, A., Sampaio, M., Jornada, J. and Radziuk, E., "SLA Design from a Business Perspective", In Proceedings of DSOM 2005.

[13] Wei, H., Frinke, D., Carter, O., et al. "Cost–Benefit Analysis for Network Intrusion Detection Systems", In Proceedings of the 28th Annual Computer Security Conference, October 2001.

[14] Boutaba, R., Xiao, J. and Aib, I., "CyberPlanner: A Comprehensive Toolkit for Network Service Providers", in Proceedings of the 11th IEEE/IFIP Network Operation and Management Symposium (NOMS 2008), Salvador de Bahia, Brazil.

[15] V. Kumar, K. Schwann, S. Iyer, Y. Chen, and A. Sahai, "A State Space Approach to SLA based Management", in *Proceedings of 11th IEEE/IFIP Network Operations and Management Symposium (NOMS 2008)*, 7-11 April 2008, Salvador da Bahia, Brazil.

[16] W. van Grembergen, "The Balanced Scorecard and IT Governance", In Proc. Int. Conference on Information Resources Management Association (*IRMA 2000*), Anchorage, Alaska, USA, 2000.

[17] Diao, Y., Keller, A., Parekh, S., Marinov, V. "Predicting Labor Cost through IT Management Complexity Metrics" in Proceedings of the 10th IEEE/IFIP Symposium on Integrated Management (IM 2007), Munich. Germany.

[18] Diao, Y., Bhattacharya, K., "Estimating Business Value of IT Services through Process Complexity Analysis", in Proceedings of the 11th IEEE/IFIP Network Operation and Management Symposium (NOMS 2008), Salvador de Bahia, Brazil.

[19] G. Barash, C. Bartolini, Liya Wu, "Measuring and Improving the Performance of an IT Support Organization in Managing Service Incidents", in proc. 2nd IEEE Workshop on Business-driven IT Management (BDIM 2007), Munich, Germany, 2007.