A Cloud-based Solution for the Performance Improvement of IT Support Organizations

Claudio Bartolini
HP Labs
Hewlett Packard
Palo Alto, CA, USA
claudio.bartolini@hp.com

Cesare Stefanelli, Davide Targa, Mauro Tortonesi
Department of Engineering
University of Ferrara
Ferrara, Italy
{cesare.stefanelli,davide.targa,mauro.tortonesi}@unife.it

***Abstract*—IT support organizations are in charge of restoring normal operations after IT service disruptions, among other tasks. Such organizations can be complex systems, with a large network of interacting support groups subject to complex management policies. The performance assessment and optimization of IT support organizations is an extremely challenging task that requires considering organization-specific structure, behavior, and business-level objectives. This paper presents Symian-Web, a decision support tool that enables IT managers to both assess and improve IT support organization performance by using what-if scenario analysis. Symian-Web features advanced information visualization concepts and metaphors, allowing for precise and timely assessment of IT support organization performance, and facilitating their re-design. Symian-Web is realized as a cloud computing-based Web application, thus enabling the tool to take advantage of the on-demand computational capabilities provided by the cloud.**

***Keywords*—Decision support tools, Information Technology Infrastructure Library (ITIL), IT service management, incident management.**

I. INTRODUCTION

IT support organizations are the entities in charge of "… *restoring normal service operation after a disruption, as quickly as possible and with minimum impact on the business*", as defined by the IT Infrastructure Library (ITIL) [1] [2]. They represent a fundamental pillar of IT service management, as they implement the *Incident Management* function.

Enterprise-class IT support organizations are usually very large systems, that implement complex organizational, structural, and behavioral processes according to the strategic objectives defined at the business management level. IT support organizations are composed of a network of support groups that cooperate to realize service restoration. In turn, each support groups has of a set of operators with specific skills and work schedules.

Because of their complexity, the performance assessment of IT support organizations is an extremely challenging task, which requires an in-depth behavioral analysis at both the system level as well as at the single support group level. In addition, the incident management process needs setting objectives that are organization-specific and defined by the business, e.g., compliance with SLAs, minimization of service restoration costs or time, etc. These must be taken into consideration during the performance evaluation phase, which therefore cannot be conducted examining IT performance metrics alone, but must consider organization-specific business-level performance metrics.

The performance improvement of IT support organizations is an even harder task, as it involves the evaluation of possible improvements caused by realignments of the current incident management strategies, as well as by the adoption of alternative strategies. The set of possible realignment operations to consider is very large, as it includes both changes at the behavioral, such as the implementation of different policies for incident assignment and prioritization, as well as at the structural level of the organization, such as support group merging, splitting, or re-staffing (increasing, cutting, or transferring operators between support groups, possibly on retraining). Notice that, in case performance realignment becomes necessary, it is essential to evaluate alternative organization structures and incident management strategies carefully before putting them in practice. In fact, the process of implementing the actual corrective measures is extremely expensive and time-consuming.

Traditional software tools provide a somewhat limited support for IT managers dealing with the performance assessment and optimization of the support organization. In fact, they usually focus on the analysis of transactional data collected from incident management logging tools, such as HP Service Manager or BMC Remedy, that do not allow to evaluate the performance of different IT support organization configurations. In addition, traditional software tools lack interfaces that enable the immediate visualization of performance data.

IT managers would benefit immensely from decision support tools that enable the analysis of incident management operations, as well as their optimization. In this context, tools based on *what-if scenario analysis*, a technique that enables the behavioral analysis of complex real life systems under alternative working conditions, are particularly interesting. In fact, by reenacting the IT support organization behavior, what-if scenario analysis tools enable to verify the effectiveness of the current structural organization, to assess how well or badly the it would perform under different loads, e.g., for stress

testing, and to evaluate possible alternative configurations when the business-level objectives are not met.

This paper presents Symian-Web, a what-if scenario analysis based decision support tool for the performance and business impact analysis, and the assisted re-design of IT support organizations. Symian-Web represents a mature and comprehensive solution, that builds on top of the results in IT support organization modeling delivered by the Symian research project [3] [4]. Symian-Web complements the accurate modeling functions delivered by Symian with a set of tools for model parameter inference, with support for a well-defined performance optimization workflow, with an innovative Web-based interface for high density information visualization, and with a cloud-based solution for the execution of the computationally expensive what-if-scenario simulations. Symian-Web represents a first significant step to push IT support organization performance optimization tools based on what-if scenario out of their current confinement in experimental research, and towards their adoption from the industry.

II. Related Work

The present work belongs to the category of research approaches that aim at improving business processes through the collection of metrics and the inferences over the acquired data [7] or the exploitation of simulation methods [8]. In this area of research, most studies propose techniques designed for business processes with complex descriptions, consisting of many well-defined steps, each with few alternative paths. Incident management, instead, is a relatively simple process taking place in a complex organization, characterized by a few simple steps with a large number of outcomes. This makes common business process management techniques either overkill or unapplicable to the incident management process.

Other research efforts have focused on optimization of IT-specific business process. Among these, a particularly interesting series of studies is Diao et al.'s, which estimate labor cost and business value of IT services analyzing process complexity [9] [10]. The present work follows a significantly different approach, and tries to improve IT support organizations via decision support and simulation techniques.

Among other research efforts in modeling IT support organizations, a particularly interesting one is Shao et al.'s EasyTicket [11] [12], a system that tries to optimize ticket routing in IT support organizations with machine learning techniques. EasyTicket is based on a queuing network model similar to Symian-Web's, but assumes that each type of incident has a specific resolver support group and tries to route incidents to their resolver group as quickly as possible. In our experience, this assumption does not always hold in practice, as usually in real life IT support organizations several support group need to cooperate in order to restore major service disruptions. The Symian-Web approach, instead, does not rely on any restrictive assumption, thus enabling a comprehensive analysis of the organization performance which considers both routing effectiveness and efficiency within support groups.

The Symian-Web interface represents IT support organizations as a network of support groups, leveraging in part on visualization concepts and methods proposed in social network analysis. Visual representation of social networks is considered essential to understand the network data and to convey the result of the analysis [13]. Symian-Web applies some social network representation techniques to the IT support organization performance optimization domain, integrating and extending them with purposely developed information visualization concepts and tools.

III. Modeling Real Life IT Support Organizations' Behavior

Symian-Web leverages on the extensive work on IT support organization modeling realized in the context of the Symian research project [3] [4]. Symian adopts a what-if scenario analysis approach based on simulation, a technique that enables the behavioral analysis of complex real life systems under alternative working conditions.

Symian allows creating realistic models of real life IT support organizations and using them to understand the impact of IT processes, strategies and tactics on the organizations' performance. Symian exploits a discrete event simulator to reproduce in detail the behavior of IT support organizations and to evaluate their performance in managing incidents. The simulation approach is particularly appropriate given that the scale and the complexity of real life organizations make it extremely difficult to devise an analytical model.

To reenact the behavior of IT support organizations through what-if scenario analysis, Symian leverages on a model based on open queuing networks and multi-server first-come-first-served queues. Markovian models such as open queuing networks are particularly attractive for the reenactment of IT support organizations, as they represent a very good tradeoff between model complexity and model accuracy. In fact, open queueing networks it can easily measure IT support organization dynamics in terms of throughput, queue lengths, response times, and utilization, both at the system level and at the single support group level. Our model has proved capable of reenacting system-wide behavior of IT support organization very well, accurately capturing end-to-end metrics such as the mean time to incident resolution and the average number of support groups visited by the incidents.

TABLE I. EXAMPLE OF TRANSACTIONAL LOGS.

Ticket ID	Sequence	Support Group ID	Service Duration (minutes)	Service Start Time
"Ticket A"	1	"SG 7"	930	12/15/2010 2.45 PM
"Ticket A"	2	"SG 53"	177	12/16/2010 7.15 AM
"Ticket A"	3	"SG 7"	34	12/16/2010 10.12 AM
"Ticket A"	4	"SG 2"	221	12/16/2010 10.46 AM
"Ticket B"	1	"SG 5"	25	1/2/2011 8.55 AM
"Ticket B"	2	"SG 19"	978	1/2/2011 9.20 AM
…	…	…	…	…

The model of IT support organizations implemented by Symian is very sophisticated and highly configurable. As a result, creating a new organization model requires a very good knowledge both of the open queuing network model realized in Symian and of the processes in place within the real life IT support organization. This challenging task is likely to be carried out by specialized personnel.

Symian also provides several purposely-designed tools that assist the user through the model parameter inference process, which is often very complicated. These tools leverage on techniques designed to build accurate models of real life IT support organizations from transactional logs, that we developed from our experience with the statistical analysis of incident transactions in enterprise-class organizations [5].

Transactional logs, as provided by incident tracking software suites such as HP IT Analytics, are the most likely information source when building a model of a real life IT support group. Unfortunately, they typically contain a very small amount of information. Usually, they do not include any information about the root cause of incidents, that would allow to identify duplicated incidents and remove their records from the dataset. Often valuable incident attributes such as their description, classification, or prioritization are also missing from transactional logs. Table I shows an example of transactional logs from a real life IT support organization. (For privacy concerns, the information provided in Table I is obfuscated.)

After the parameter inference process, Symian-Web lets the user check whether the model just created reproduces the IT support organization baseline behavior with an acceptable degree fidelity. To this end, Symian-Web enables to feed the support organization model with historical incident arrival traces and to compare its performance with that of the corresponding real life IT support organization under the same conditions.

IV. SYMIAN-WEB

Symian-Web integrates and extends the well-tested state-of-the-art IT support organization model provided by Symian to realize a comprehensive decision support tool.

In fact, Symian-Web implements a streamlined performance optimization workflow that guides IT managers in every step of the optimization process, from the IT support organization model creation to its assisted re-design. In addition, Symian-Web allows the definition of custom performance analysis functions, an essential feature to enable the optimization according to business-driven criteria.

Symian-Web also proposes an innovative interface that displays IT support organization structural and performance information simultaneously, thus enabling a deep insight into the organization dynamics. The intuitive and easy to use Symian-Web interface features high-density visualization mechanisms that provide immediately accessible information at different aggregation levels. At the same time, Symian-Web enables users to change the IT support organization configuration directly from the performance visualization interface, guiding them through the available set of reconfiguration options.

Symian-Web is realized as a cloud computing-based Web application, adopting a Software-as-a-Service (SaaS) approach. Cloud computing allows Symian-Web to leverage on elastic and on-demand computation resources to carry out simulations of IT support organization behavior. The SaaS approach also facilitates both the distribution and the management of updates compared with traditional desktop GUI applications. In addition, Web access facilitates demonstrations to researchers and practitioners working on the IT support organization performance optimization, that can thus evaluate our software at their convenience and provide valuable feedback. Finally, the "perpetual beta" software provisioning model enabled – and frequently adopted – by SaaS applications represents a perfect match for research projects such as Symian-Web.

V. THE INTERACTIVE IT SUPPORT ORGANIZATION PERFORMANCE OPTIMIZATION PROCESS

After the parameter identification phase, Symian-Web users can begin playing out what-if scenario analysis with the inferred IT support organization model. In Symian-Web each support organization model can have many *configurations*, representing different structures, e.g., merged, split, or restaffed support groups, or working conditions, e.g., a higher incident load for stress-testing purposes.

Symian-Web enables users to compare the performance of different configurations for the same IT support organization model, and verify which one represents the best alignment to the business objectives. This effectively enables an iterative performance optimization process, in which users can incrementally specify the set of changes to apply to the current organization model in order to define an alternative configuration that will be tested on a set of performance metrics.

When a new IT support organization model is created, it is in a *default configuration*. New configurations can be created by cloning an existing one and changing the desired

TABLE II. VISUAL METAPHORS ADOPTED BY SYMIAN-WEB TO REPRESENT IT SUPPORT ORGANIZATION PERFORMANCE.

Visual Metaphor	Represents	Rationale
Proximity between nodes	*Strength of interactions between the corresponding support groups*	Support groups that communicate often are depicted next to each other, enabling users to immediately detect the support groups that are good candidates for merging.
Node size	*Amount of work performed by the corresponding support grop*	Size of the node is proportional to the number of tickets processed, enabling users to immediately detect the most important support groups.
Edge size	*Flow of tickets between the corresponding support groups*	Edges come to represent "information highways", thus enabling users to immediately detect the most important paths that incidents take through the organization, as thicker edges indicate a larger flow of incidents.
Node border	*Number of tickets closed by the corresponding support group*	This enables user to immediately see which support groups are the most effective in finalizing the service restoration process.
Node color	*Support level*	When information about the level of support (Help Desk, first line of support, second line of support, etc…) of support groups is available, coloring is used to represent it.

parameters. Most users will find this differential specification – that is, the definition of a new configuration with respect to a preexisting one – the most convenient way to create a new IT support organization configuration. However, experienced users may instead prefer to take advantage of the secondary, lower-level interface provided by the tool. In fact, Symian-Web also allows to create new configurations by uploading a file containing a full IT support organization configuration specification.

When created, a model configuration is in the "Configurable" state. In this state, users cannot immediately access performance information, but they can perform any reconfiguration they want on the IT support organization model configuration. Any attempt by the user to visualize performance result will automatically trigger a simulation run – unless performance data is already available in the tool cache – and a switch of the IT support organization model configuration state to "Running". In the "Running" state, users cannot perform any reconfiguration of the IT support organization model configuration. Visualization of performance results is also prevented. However, users are not simply put on hold, as they can copy the current IT support organization model configuration and work on the cloned configuration. When the simulations terminate, the IT support organization model configuration state becomes "Assessed", and performance information is finally accessible to the user. In this state, any IT support organization model configuration reconfiguration will bring it back to the "Configurable" state.

Once an IT support organization configuration is finalized, Symian-Web automatically re-enacts it through a series of simulations and then lets the user assess its performance through a purposely-developed high density information visualization interface.

Simulations can take a rather long time to complete, in the order of a few minutes for large IT support organization. And, in order to obtain statistically significant results, Symian-Web performs several different simulations (by default 20, but configurable) every time an IT support organization configuration is modified. As a result, Symian-Web runs simulations asynchronously with respect to the user interactions. In fact, after they trigger a new simulation run, users cannot be expected to wait for the job to complete without refreshing the page or hitting the "Back" button of the browser – a potentially harmful operation that might leave the system in an inconsistent state.

VI. VISUALLY AIDED PERFORMANCE ASSESSMENT AND MODEL RECONFIGURATION

To analyze the performance of the reenacted IT support organizations, Symian-Web users can exploit the high-density interactive information visualization interface provided by the tool, that enables the immediate and comprehensive visualization of performance information.

We designed the Symian-Web interface according to our experience with IT support organization performance optimization as well as on our observations on how IT managers process performance-related information. In fact, in our experience performance bottleneck spotting and removal often represents a very effective strategy to optimize the performance of IT support organization models. In addition, we noticed that IT managers tend to attribute a significantly higher importance to high-level aggregated performance information than to lower-level details. These consideration suggested the development of an interface that provides several levels of details as well as a deeper focus on a selected portion of the organization.

The main view of Symian-Web is a map of the support groups of the IT support organization. Symian-Web adopts a graph-based representation to provide simultaneously structural and performance-related information about the IT support organization. More specifically, each node in the graph represents a support group of the IT support organization and each edge represents the presence of ticket escalations between the corresponding support groups. Fig. 1 shows the Symian-Web graph-based interface for the model of a real life enterprise-class IT support organization.

The graph-based interface leverages on several *visual metaphors*, which we introduced in [6], to facilitate digesting information, and proposes a streamlined interface designed to keep all the functions always easily accessible. The visual metaphors on which Symian-Web leverages are summarized in Table II.

The Symian-Web interface enables interactive data visualization. Nodes can be moved (by drag and drop), the view on the graph can be zoomed in or out, etc. Users can select both the desired detail level and focus point. The interface also presents additional contextual information, such as secondary attributes of support groups and their dynamics, at the bottom of the graph-based IT support organization representation.

The interface also enables the visual reconfiguration of the IT support organization mode. Symian-Web provides several options to optimize the performance of IT support organizations. Some of the operations available to IT managers, such as support group removal, support group creation, merging of two support groups, and splitting of a support group, might have a major impact on the IT support organization model. When two (or more) workgroups are highlighted, by control-clicking on their representative nodes, the Symian-Web interface enables the merge operation. Splitting support groups is done by highlighting a group and selecting the split option. Initially, the newly created group is supposed to receive half of the ticket traffic of the original workgroup, and have a fan-out that is identical to the original group. However, a slider-like control can be used to determine the share of the tickets that newly created groups are expected to receive.

When the user finalizes the changes to the IT support organization model, the Symian-Web interface triggers a new simulation run to assess the performance of the new model configuration, displays an out-of-date information warning message, and prevents further changes to the model configuration until the simulation outcome is available. At the end of the simulation, the Symian-Web interface updates the graph-based representation with data collected from the simulation.

VII. Custom Performance Analysis

At the end of each simulation, Symian-Web calculates both Key Performance Indicators (KPIs) and the IT metrics that are of crucial importance in understanding the IT support organization performance. To this end, Symian-Web leverages on a library of predefined metrics as well as on custom metrics defined through a dedicated scripting language. This enables to fully customize the process of performance assessment according to the specific needs and characteristics of the IT support organization – an essential feature in order to enable business impact-driven performance analysis.

Users can then analyze the collected performance data using the Symian-Web reporting functions. Symian-Web enables the visualization of available data on the dynamics of the IT support organization, as well as of each support group, according to different dimensions. For example, a user may be interested in a specific time period, incident category or severity. Symian-Web supports several plot types, as well as tabular representations, for performance information visualization. To present hierarchically organized data in a compact and immediately understandable way, Symian-Web also leverages on interactive high-density visualization methods such as Treemaps, Sunburst, and Icicle diagrams.

The mapping of visual features to metrics for support group dynamics in the Symian-Web interface is also highly customizable. For example a user may want the dimension of the node to indicate the staffing level of the workgroup, or the coloring to indicate the geography that the workgroup belongs to. Another possibility is to represent the input and output, i.e., incoming and resolved incidents, of the IT support organizations as connections with two virtual support groups,

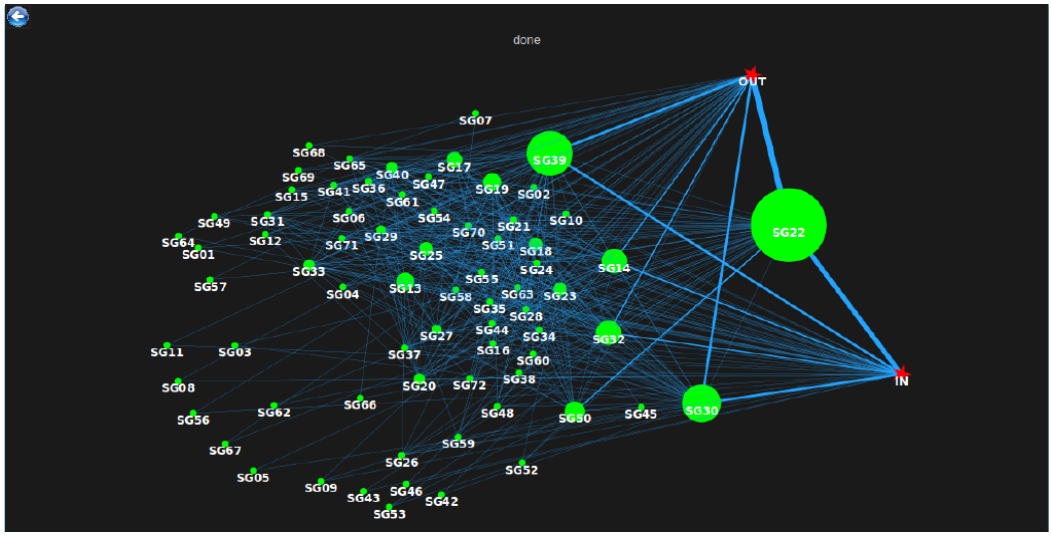

Figure 1. The Symian-Web graph-based representation of IT support organization.

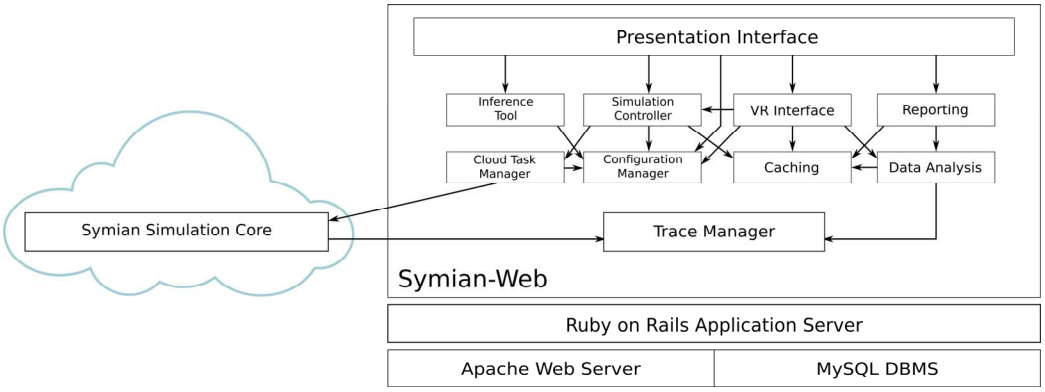

Figure 2. The Symian-Web Architecture.

respectively dubbed "IN" and "OUT", as shown in Fig. 1. Finally, node vicinity (the map layout) could also be customized to represent other features. For example, the nodes could be layered as to represent the support level hierarchy or even drawn on a geographic map.

VIII. Example of What-If Scenario Analysis for Reconfiguration of IT Support Organization

In this section, we provide an example of the kind of analysis that the Symian-Web tool is able to provide, by describing some reconfiguration experiments that we run on real life data provided to us by the Outsourcing Services Division of HP. The data we used is taken from transactional logs of 34 support groups of HP Outsourcing dedicated to a financial services customer, plus 38 more groups that have shared responsibilities across multiple enterprise customers. We were able to obtain database logs of incidents for a 377-day period, consisting of data for about 25,000 incidents. For each incident, the data carried transactional information about the arrival and departure times at each visited support group. We used the data to derive a model of the organization following the method described in section IV above, and fully explained in [4].

At the end of the configuration process, we ran a first series of simulations to get evidence for the accuracy of the model we built for the IT support organization. Fig. 1 shows the performance visualization interface provided by Symian-Web for the IT support organization model.

Following the model construction and validation process, we then used Symian-Web to optimize the IT support organization model with respect to the *service disruption time* parameter, with the constraint of preserving the current number of operators in each of the support groups. As a result, the objectives of the performance improvement process are the maximization of the Mean Incidents Closed Daily (MICD) metric, as well as the minimization of the Mean Time To incident Resolution (MTTR) metric.

Using the Symian-Web performance information interface, we immediately identified support group 22 (the rightmost node in Fig. 1) as the group dealing with the largest number of tickets, followed by support groups 39 and 30. We then considered support group 22 as the most critical performance bottleneck and tried to change the IT support organization model configuration in order to mitigate this problem.

As an example of assessing performance improvement based on organizational redesign using Symian-Web, we increased the operator effectiveness at one of the support groups marked as a performance bottleneck by 2%, emulating an improvement in operator performance that could be obtained in real life by re-training technicians. We then launched a new series of simulations to assess the impact of this change on the organization performance. Table III provides a comparison of the MICD and MTTR performance metrics before and after the optimization. The results of the second series of simulations demonstrated that the bottleneck location and removal is an effective method to improve the whole system performance. In particular, the IT support organization exhibited a 1.00% improvement of the MICD and a 4.91% decrease of the MTTR.

TABLE III. Performance Metrics Comparison from the First and Second Series of Simulations (with a 95% Confidence Interval).

	Before optimization	**After optimization**
MICD	140.25±0.34	141.17±0.30
MTTR (in seconds)	10010027±18315	951860±21949

IX. Architecture and Implementation

The Symian-Web architecture follows the Model-View-Controller pattern, that enables separation of concerns with regards to the 3 essential functions of IT support organization modeling, performance information representation, and response to user interactions. Symian-Web adopts a layered architecture, as depicted in Fig. 2 which highlights the tool's main components. The top layer takes care of common presentation functions, with the Presentation Interface module that interfaces with lower-layer components to provide access to all the Symian-Web functions.

The middle layer contains components that implement the main Symian-Web functions. The VR Interface represents the

main interface to the functions provided by Symian-Web, as described in Section VI. The Reporting module provides reporting functions on IT support organization configurations' performance data. The Simulation Controller keeps track of each IT support organization configuration state and manages simulation runs in the cloud. The Inference Tool module provides functions, such as statistical data analysis ones, that support and guide users in the creation of IT support organization models.

The third layer provides cloud-based support for the execution of simulations as well as common functions for the upper layer components. The Cloud Task Manager component triggers and supervises the execution of simulations in the cloud, dispatching them to the cloud for execution and monitoring them. The Configuration Manager provides functions to store, retrieve and compare IT support organization configurations and the related simulation parameters. The Caching module provides caching functions that allow to save simulation outcomes and post-processed performance data in order to avoid the useless re-execution of expensive computations. The Data Analysis module performs the post-processing of the results collected from simulation execution, leveraging on predefined as well as custom user-provided metrics to evaluate the IT support organization performance.

The lowest layer takes care of simulation execution tasks. The Symian Simulation Core is the component that actually executes IT support organization configuration simulations, and thus runs within the cloud. Symian Simulation Core provides simulation outcomes to the Trace Manager component, which logs all the information collected during the execution of the simulation. This raw information will then be later post-processed by the Data Analysis module.

The Web approach also enables Symian-Web to expose a ReST-based API [14] [15] that makes the tool remotely controllable. The ReST-based API allows every operation accessible through the Web interface, such as inserting new IT support organization models, adding a configuration to an already existing IT support organization, launching simulations, and retrieving simulation result data. This enables the integration of Symian-Web with other tools, e.g., for the post-processing of simulation outcomes or the automation of long and/or common procedures.

Symian-Web is implemented as a Ruby on Rails application. Ruby on Rails is a framework for the development of Web 2.0 applications that allows the rapid development of new features, minimizing their time-to-deployment [16] [18]. Symian-Web uses the JavaScript InfoVis Toolkit (*http://thejit.org/*) to realize the graph- and tree-based IT support organization representations implemented by the VR interface. Symian-Web also integrates with the R statistics framework (*http://www.r-project.org/*) to realize the complex statistical data analysis functions provided by the Inference Tool component and to produce high quality plots of the simulation results.

X. Conclusions and Future Work

The main contribution of Symian-Web is that it provides a comprehensive framework for the performance optimization of IT support organizations. The tool leverages on state-of-the-art what-if scenario analysis functions developed in the context of the Symian research project to reenact the organization behavior under different configurations or working conditions. Symian-Web complements those IT support organization modeling functions with specific support for the realistic model, the implementation of a streamlined performance optimization workflow, an innovative Web-based interface for information visualization, and solutions for the execution of the computationally expensive simulations in the cloud.

In our evaluation of Symian-Web, we observed that the what-if scenario analysis technique coupled with the high density information visualization functions provided by the tool is very effective in detecting performance bottlenecks within IT support organizations.

We are currently investigating other higher-density information visualization methods for IT support organization performance data. We are also considering further evaluations of Symian-Web, involving interested researchers and practitioners as well as IT managers and collecting their feedback to improve the feature set and usability of the tool.

References

[1] Office of Government Commerce, "Service Strategy Book", The Stationery Office, Norwich , 2007.

[2] Office of Government Commerce, "Service Operation Book", The Stationery Office, Norwich, 2007.

[3] C. Bartolini, C. Stefanelli, 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, Pythagorion, Samos Island, Greece.

[4] C. Bartolini, C.Stefanelli, M. Tortonesi, "SYMIAN: Analysis and Performance Improvement of the IT Incident Management Process", *IEEE Transactions on Network and Service Management*, Vol. 7, No. 3, pp. 132-144, September 2010.

[5] C. Bartolini, C. Stefanelli, M. Tortonesi, "Modeling IT Support Organizations from Transactional Logs", in *Proceedings of the 12th IEEE/IFIP Network Operations and Management Symposium (NOMS 2010)*, pp. 256-263, 19-23 April 2010, Osaka, Japan.

[6] C. Bartolini, C. Stefanelli, M. Tortonesi, D. Targa, "A Web-based What-If Scenario Analysis Tool for Performance Improvement of IT Support Organizations" (short paper), accepted for publication in *Proceedings of 7th International Conference on Network and Service Management (CNSM 2011)*, 24-28 October 2011, Paris, France.

[7] F. Casati, M. Castellanos, U. Dayal, M.C. Shan, "A Metric Definition, Computation, and Reporting Model for Business Operation Analysis", in *Proc. of Advances in Database Technology - EDBT 2006, 10th International Conference on Extending Database Technology*, Munich, Germany, March 26-31, 2006.

[8] K. Tumay, "Business process simulation", in *Proc. Winter Simulation Conference 1995*, 3-6 December 1995.

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

[10] Y. Diao, K. Bhattacharya, "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.

[11] Q. Shao, Y. Chen, S. Tao, X. Yan, N. Anerousis, "Efficient Ticket Routing by Resolution Sequence Mining", in *Proceedings of 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD'08)*, Las Vegas, NV, USA, 24-27 August 2008.

[12] Q. Shao, Y. Chen, S. Tao, X. Yan, N. Anerousis, "EasyTicket: A Ticket Routing Recommendation Engine for Enterprise Problem Resolution", in *Proceedings of 34th International Conference on Very Large Data Bases (VLDB'08)*, Auckland, New Zealand, 23-28 August 2008.

[13] L. Freeman, "Visualizing Social Networks", *Journal of Social Structure*, Vol. 1, No. 1, 2000.

[14] R. Fielding, R. Taylor, "Principled Design of the Modern Web Architecture", *ACM Transactions on Internet Technology*, Vol. 2, No. 2, pp. 115–150, May 2002.

[15] L. Richardson, S. Ruby, "RESTful Web Services", O'Reilly, 2007.

[16] I. Vosloo, D. Kourie, "Server-Centric Web Frameworks: An Overview", *ACM Computing Surveys*, Vol. 40, No. 2, Article 4, April 2008.

[17] M. Bachle, P. Kirchberg, "Ruby on Rails", *IEEE Software*, Vol. 24, No. 6, pp. 105-108, 2007.