

IT Incident Management as a Collaborative Process: A Visualization Tool Inspired to Social Networks

Claudio Bartolini

HP Laboratories
1501 Page Mill Rd, Palo Alto, CA 94304
claudio.bartolini@hp.com

Abstract. Dynamics of IT support organizations can be quite complex to analyze. We introduce a prototype visualization tool that exploits similarities of support organizations with social networks. Our tool borrows techniques from social network visualization and analysis to help IT manager better understand dynamics of their organization in the incident management process.

Keywords: IT service management, IT infrastructure library, incident management, social networks, visualization.

1 Introduction

The IT Infrastructure Library (ITIL [1]) is a comprehensive set of concepts and techniques for managing IT infrastructure, development, and operations. Developed by the UK Office of Government Commerce (OGC), ITIL is today the de facto best practice standard for IT service management. Among the processes that ITIL defines, *Incident Management* is the process through which IT support organizations manage to restore normal service operation after a disruption, as quickly as possible and with minimum impact on the business.

IT support organizations are structured as complex networks of support groups, each comprising of a set of skilled technicians. Incident management can therefore be seen as a very complex collaborative process. With the complexity of IT support organizations and staff working around the clock in disparate geographies, it is hard to understand the dynamic of the organization.

In this paper, we present ITSupportster, a visualization tool that takes inspiration from graph-like representation of social networks. The remainder of this paper is structured as follows: in section 2 we present an abstraction of the IT incident management process, which we used as a basis to build the model behind our visualization tool; in section 3 we describe our visualization tool, ITSupportster; in section 4 we present empirical and anecdotal evidence of the usefulness of our tool; we review related work in section 5 and conclude in section 6 mentioning the next steps in continuing development of our prototype tool.

2 The IT Incident Management Process

An IT support organization (fig. 1) consists of a network of support groups, each comprising of a set of skilled technicians, with their work schedule. Support groups are divided into support levels (usually three to five). The helpdesk is assigned level zero, support groups at lower levels may be able to deal with generic issues such as “user forgot password”, while groups at higher levels are more specialized and able to deal with harder tasks. Support groups are further specialized by category of incidents that they treat (network, server, etc...) and divided into geographies, to ensure a more prompt handling of incidents especially at lower support levels.

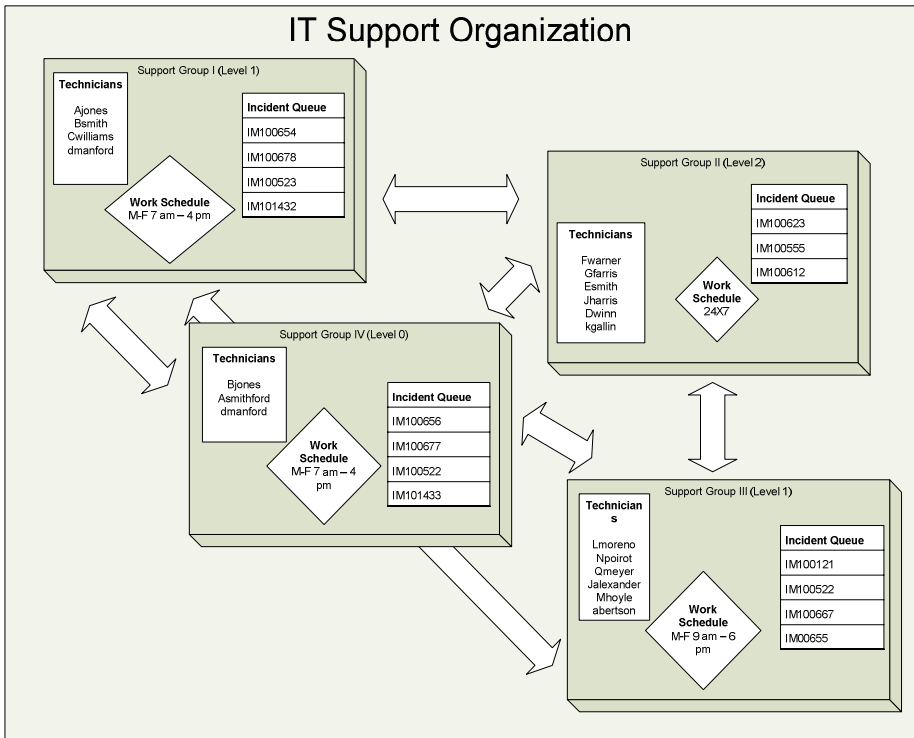


Fig. 1. Conceptual model of the IT support organization for incident management

When a customer experiences a disruption in an IT service, he or she interfaces with a help desk, which opens an incident, also sometimes called troubleticket (simply ticket in the rest of our paper). The incidents are worked on by different support groups throughout their lifetime, as technicians in a given support group can either resolve the incident or pass it on to a different support group (usually escalating to a higher level of support). We model the support organization as a queue system, each support group having a queue of incidents that it is working on at a given time. This process is exemplified in fig. 2.

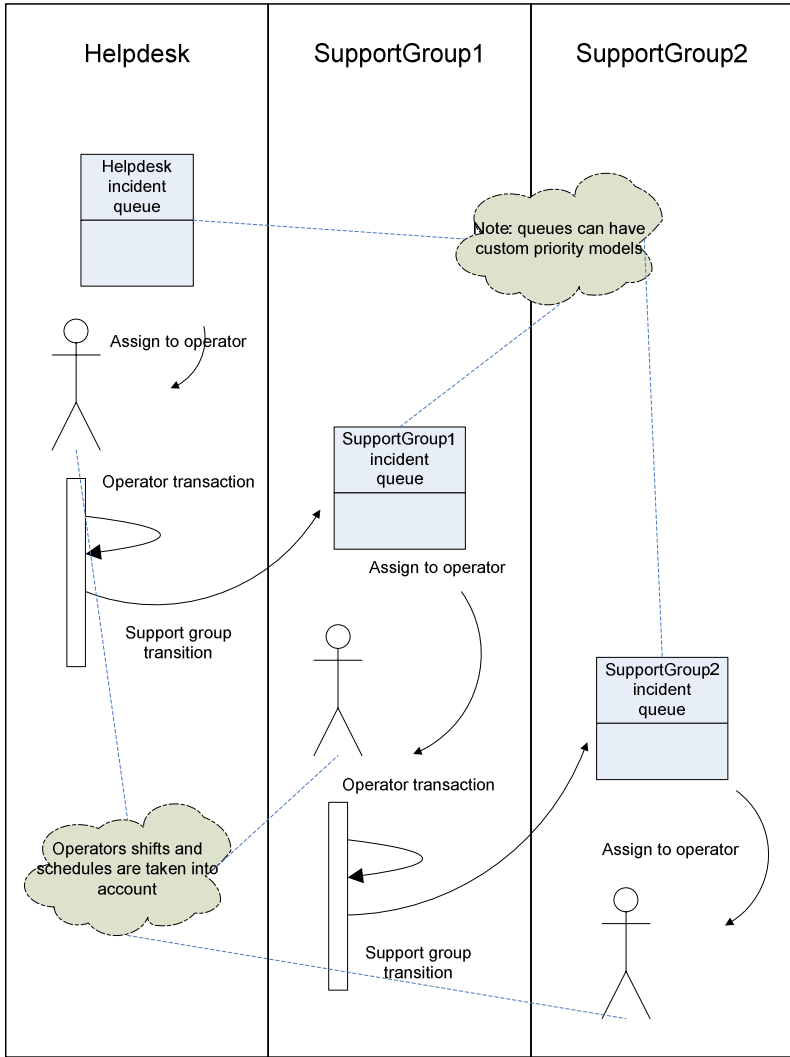


Fig. 2. The incident management process

The incident lifecycle (fig. 3) is described as the progression of an incident from its initial occurrence to detection, then to the diagnosis of the cause of failure, the repair of the issue, recovery of the component that was at fault, and finally restoration of service. Throughout the lifetime of an incident, the incident gets "opened" by the helpdesk, then assigned to a support group which is going to work on it and either mark it "closed" or "reassign" it to a different support group. At each of these states, the incident record is updated with the pertinent information, such as the troubleshooting that has taken place, which team is responsible for action, which engineer on the team, and what the current status of the incident is. Figure 3

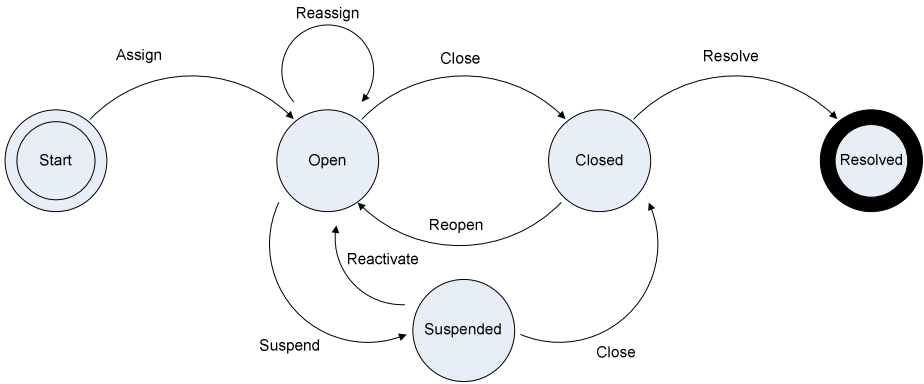


Fig. 3. Incident lifecycle

demonstrates the progression between these different states: An incident is opened when a call comes in to the call center after an issue is detected.

Once the incident is assigned to the appropriate team that will handle the issue, it is updated numerous times to reflect the progression of troubleshooting that the engineers attempt. If, for some reason, the end-users requests that the engineer stops working on the issue, it is placed in a “suspended” state so as not to incur SLO penalty while not being worked on. Once the disruption is repaired, the incident is placed in “closed” state until the end-user confirms that service has been restored and all is well. Once the restoration is confirmed, the incident is “resolved” and its lifecycle ends.

3 ITSupportster

ITSupportster is a visualization tool that provides insight into the structure and the dynamics of an IT support organization.

The main view of ITSupportster is a map of the workgroups that compose the IT support organization, as shown in figure 4. Workgroups are laid out on the map according to inter-group communication dynamics: workgroups that communicate often – i.e. redirect tickets from one another – are near to each other on the map. This is similar to the concept of social networks where entities that are marked as “friend” are drawn near to one another. All of the data necessary to populate the ITsupportster social network map is collected through off-the-shelf IT help desk software products. To develop our prototype, we gathered data from installations of HP Service Manager [2].

Various pictorial features on the map are used to convey information. A prominent feature of the visualization is that the size of the node is proportional to the number of tickets that the workgroups process. In this way it is immediate to get an impression of what workgroups see the most traffic. The thickness of the directed edges connecting nodes is used to represent the ticket throughput between the workgroups. A thicker edge represents a higher number of tickets being passed from one group to the next. In this edges come to represent “information highways”. The thickness of

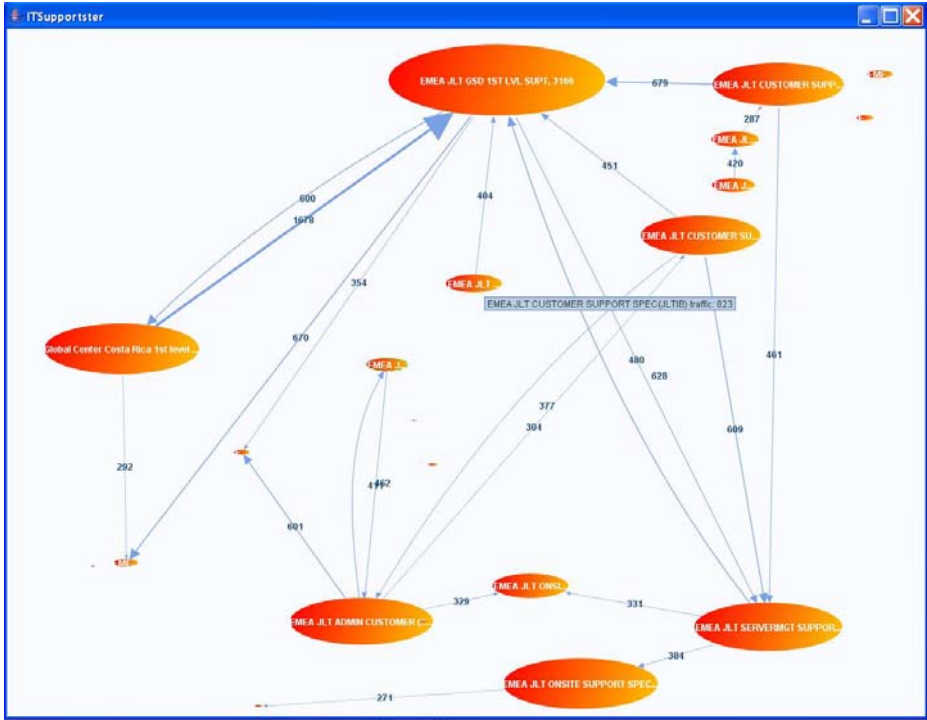


Fig. 4. ITSupportster main view

the node border is proportional to the number of tickets resolved by the support group represented by the node. The coloring is used to represent the level of support (helpdesk, first line of support, etc.)

It has to be noted that the mapping of visual features to metrics for support group dynamics is 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. Finally, node vicinity (the map layout) could also be customized to represent other features. For example, the nodes could be layered as to represent to support level hierarchy or even drawn on a geographic map.

Given the elevated number of nodes on the graph (the number of workgroups in an IT large support organization can be up to thousands), the map is represented through a hyperbolic graph [3]. The main feature of a hyperbolic graph is that the region of the map that is in focus is rendered with greater detail, while still giving a picture comprising of the whole map. Zoom-in and zoom-out are supported.

Because of the extremely high density of information that it conveys, the main ITSupportster map is the ideal starting point for IT performance analysis, business impact analysis and assisted optimization (re-design) of the IT support organization.

Highlighting a node representing a workgroup on the map, a side pane (workgroup view) shows structural information about the workgroup (name, manager, contacts, location, technicians, etc...), and collates information useful for IT performance

analysis that is collected through a separate decision support module, SYMIAN described in more detail in [4, 5]. The architecture of ITSupportster follows the Model-View-Controller architectural pattern. The social network map that is used as a view in the IT performance analysis and business impact analysis modalities, is also used as controller in the guided optimization (or assisted organization design) that is available through the sister tool SYMIAN.

4 Validation of the Prototype

We built several visualization examples from real life data derived from installations of HP Service Manager in an outsourcing context. When we shared the examples with the IT managers, they immediately recognized their organization, and in more than one case were able to suspect workgroups to be bottlenecks in the network. In one example, one major flow identified by ITSupportster as a possible red herring, was already suspected to be an issue. However they had no quantitative data to support this, and the diagram and associated metrics helped to focus the dialog between the two workgroups involved, contributing to the issue resolution.

5 Related Work

Our work on visualization of incident management dynamics can be seen as part of a recently emerged research area that focuses on the two IT management dimensions of people and processes. The interest for this topic arose as researchers started analyzing the relationships between people, processes and technological optimization and the impact of automation and process complexity on labor cost. As a representative example, we cite Diao et al.'s recent research effort addressing the very important question of when does it make sense to automate processes based on metrics of process complexity [6, 7]. The main difference of our approach and theirs is that make similar modeling considerations to achieve significant improvements in the performance of the organization through visualization, decision support and simulation techniques.

The analysis of the incident management process and the IT support organization model presented in this paper shares common aspects with what we described in [8] and then again in [4]. In this paper, however, we build on those abstractions to build a tool that is able to visualize the dynamics between the various support groups that we began to describe there.

Social network analysis has been used as an analysis technique applied to various domains ranging from information science, organizational behavior, economics, and other. Social network analysis explores the relationships and ties between actors within the network, rather than the individual aspects of each actor. Visual representation of social networks is deemed important to understand the network data and convey the result of the analysis. See for example Freeman's work [9]. Our ITSupportster tool can be seen as an application of these ideas to the IT support organization seen as social network of support workgroups.

6 Conclusion and Future Work

Dynamics of IT support organizations can be quite complex to analyze. We presented ITSupportster, a prototype visualization tool that exploits similarities of support organizations with social networks and borrows techniques from social network visualization and analysis to help IT manager better understand dynamics of their organization in the incident management process.

Our next steps will be to closely couple the ITSupportster visualization prototype with a decision support tool that we have developed for optimizing the performance of IT support organizations through simulation techniques (SYMIAN, described in [4]).

References

1. IT Infrastructure Library, ITIL Service Delivery and ITIL Service Support, OGC, UK (2003)
2. HP ServiceManager, <http://www.hp.com/software>
3. Herman, I., Melançon, G., Marshall, M.S.: Graph Visualization and Navigation in Information Visualization: A Survey. *IEEE Transactions on Visualization and Computer Graphics* (2000)
4. Bartolini, C., Stefanelli, C., Tortonesi, M.: SYMIAN: a Simulation Tool for the Optimization of the IT Incident Management Process. In: De Turck, F., Kellerer, W., Kormentzas, G. (eds.) *DSOM 2008*. LNCS, vol. 5273. Springer, Heidelberg (2008)
5. Bartolini, C., Stefanelli, C., Tortonesi, M.: Business-impact analysis and simulation of critical incidents in IT service management. In: *Proceedings of IEEE Integrated Management Conference IM* (2009)
6. 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 (2007)
7. 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 (2008)
8. Barash, G., Bartolini, C., Wu, L.: 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)
9. Freeman, L.C.: Visualizing Social Networks. *Journal of Social Structure*, JoSS 1(1) (2000)