Using "Rapid Experimentation" to Inform Customer Service Experience Design

Soni Meckem
Director, Service Transformation
Cisco Systems, Inc.
170 West Tasman Drive,
San Jose, CA 95134 USA
smeckem@cisco.com

Jennifer Lee Carlson
Senior Consultant for Cisco
Tec-Ed, Inc.
4300 Varsity Drive, Suite A,
Ann Arbor, MI 48108 USA
jennifer@teced.com

CHI 2010, April 10 – 15, 2010, Atlanta, GA, USA
ACM 978-1-60558-930-5/10/04.

Abstract

This case study describes how Cisco followed a "Rapid Experimentation" methodology in conducting iterative, high velocity pilot studies to inform a large global customer service experience design project. The research findings described in this case study informed the design of a better mechanism for customers to select their expected outcomes, so Cisco can provide a personalized service experience. This improved accuracy moves us closer to our goal of eliminating at least 5% of all re-routing of service requests. In addition, customer satisfaction improves as we approach our target of reducing average Time-To-Resolution by at least 5%, which also saves on the Cost-Per-Call for Cisco.

The case study explains how these studies improved the direction of the design concept and narrowed the research focus to answer more specific design questions. It summarizes how this approach was successfully applied in the customer service experience design situation to achieve the same experience design goal in 8 weeks, 4 weeks ahead of the 12-week schedule. We also describe lessons learned in applying the "Rapid Experimentation" methodology.

Author Keywords
Customer service experience, service design, customer support experience, Rapid Experimentation, Rational Unified Process, iterative studies, Rapid Iterative Testing and Evaluation (RITE), comparative usability study, Agile Software Development, Toyota Production System (TPS).

ACM Classification Keywords
H5.3. Information interfaces and presentation (e.g., HCI): Group and Organization Interfaces: Evaluation/methodology.

General Terms
Experimentation, Design

Introduction
Part of a Broader Project
In June of 2009, Cisco formed a new End-to-End customer service experience design [1] team (End-to-End Team) chartered to transform the experience of customers who contact Cisco to have their technical issues resolved. This End-to-End Team immediately carried out a global contextual observation study that produced more than 50 conceptual experience improvement ideas. The next step was to collect more data to refine these ideas.

The Need for Speed
To be successful, we knew it had to be rapid and take days or a few weeks—but not more than a month—to explore each conceptual experience improvement idea. Cisco's technical support services are a large contributor to overall customer satisfaction and loyalty. We sought ways to accelerate the customer service experience design process to quickly implement meaningful change.

The "Rapid Experimentation" Approach
The team's director created a "Rapid Experimentation" approach using lessons learned from applying the principles of the Toyota Production System [6], Agile Software Development [2, 5], and the Rational Unified Process [3]. The approach employs iterative and high velocity testing of an idea using low-fidelity simulation; it has some similarities to the RITE method [4]. While the End-to-End Team pursued the larger customer service experience design project, it spun off ideas for exploration to four small research teams that would use the "Rapid Experimentation" approach to conduct separate studies in parallel. The three studies described next were conducted in sequence by one of these small teams.

We decided to apply the following principles for the "Rapid Experimentation" method:

- Use small research teams, work quickly
- Consider hypotheses that are extremely narrow in scope
- Seek confidence levels that are directionally correct but not necessarily statistically significant, since their purpose is only to inform the larger study
- Use low-fidelity techniques, saving higher-fidelity techniques for the iterative approach used during build and implement

Trying the Approach
Customers had distinctively different expectations when contacting Cisco. Currently, there is a single process for requesting service, regardless of the customer's desired

outcome. We wanted to design a mechanism to capture customer-specified outcomes early in the process. This would allow us to tailor the rest of the customer experience according to customers' expected outcomes, enabling them to provide only the information relevant to solve their particular issue and giving them quicker access to the appropriate service personnel. Our study looked at improving the process of gathering accurate information about a customer's expectations. The End-to-End Team initially identified the following six customer expectation categories: restoration, Return Materials Authorization (RMA), troubleshooting, expert consultation, information, and root cause.

The series of three studies described here, conducted by one of the spun-off research teams, were to inform the design and assess the feasibility of such an approach. It was critical to the End-to-End Team's design project that the research team complete the first study in two weeks, because the result of this study would fundamentally change the entire design if the hypothesis proved not to be true.

First Case Study — Are there a manageable number of expected outcome categories?

To investigate how Cisco can better provide a differentiated experience tailored to a customer's expectations for a service request, the research team first needed to test the hypothesis that there are a manageable number of expected outcome categories.

Study Design

In order to confirm or disprove the hypothesis, we had to answer three questions:

- How well do the target customer's expectations fit into the six outcome categories?
- If there are other categories besides those six, what are they?
- Do the naming and categorization of the six categories make sense to the customers?

We conducted phone interviews with customers and in-person interviews with internal Support Engineers to collect both quantitative and qualitative data. Quantitative data included the number of times participants reported the categories as their expected outcomes and the number of additional categories participants volunteered. Qualitative data included participants' comments on the accuracy of the categorization, the predefined category names, and reasons why some outcomes do not fit into the predefined categories.

Rather than asking the customer participants to evaluate the suitability of the predefined categories using a numbering scale, we asked open-ended questions to collect participants' own words for their expectations. The six predefined categories were not mentioned or shown. If participants described their expectations using other terminology, we prompted for differences or similarities between the participant's language and the predefined name for a given category.

Customer participants were technical staff from top-tier customers who had created a service request at Cisco within the past 45 days. Cisco has customers all over the world, so our study included customer participants from different regions to ensure that it was representative of the customer base. There were 10

customer participants: 7 from the US and Canada, 1 from Europe, 1 from Asia Pacific, and 1 from emerging markets.

Cisco internal participants consisted of Support Engineers from two support organizations that provide different levels of expertise, depending on the complexity of the service request and the level of service to which customers are entitled.

Before the study sessions, we developed a study plan, recruiting materials, session materials, and a data collection template.

Execution

The study took place in July, 2009. It required approximately 120 hours of work over 14 days, including planning, material preparation, recruiting, conducting the sessions, data analysis, and reporting. The research team consisted of a full-time researcher, a part-time recruiter, and a part-time note-taker.

During the 30-minute phone interviews with each participant, the researcher asked, "What are the reasons for you to contact Cisco?" and "What are the expected outcomes?" From their answers, we captured their words describing their expectations, and asked "Is [*customer's word*] the same or different from [*the most similar predefined category name*]?" and "How are they different to you?" If the participant didn't mention any of the six predefined categories, we asked "Have you needed [*the category*] from Cisco?" and "Please tell me more."

To accelerate the schedule, all the customer interviews were performed by a team of two: moderator and note-taker. By the last day of the sessions, all data were logged into the data collection template and summary data were generated with the predefined formulas in Excel.

Six Cisco Support Engineers participated in this study. Besides asking similar questions about what their customers expected from Cisco, we also asked each participant to rank the outcome categories by frequency of requests.

Results

The study confirmed that there is a manageable number of expected outcome categories from the customers. It gave the End-to-End Team confidence to move forward with the outcome-based experience design direction.

All 10 customer participants reported that they had opened service requests for troubleshooting and RMA. Seven participants said that they had opened service requests for expert consultation. Fewer participants had opened service requests for restoration, root cause, and information. None of the customers identified new outcome categories, although several suggested different names for them (Figure 1).

Category	# of Responses	Customer's Words
RMA	10	Replace an equipment
Troubleshooting	10	Break fix; Issue (or problem) resolution; Fix an issue
Expert Consultation	7	Configuration assistance; Recommendations; Tell me how to…
Restoration	3	Restore a functionality; Break fix
Root Cause	1	Tell me why…
Information	1	Help me find the answer/information in cisco.com

Figure 1: Number of Customer Responses for each Category

Four of the 6 Cisco Support Engineers ranked troubleshooting and 2 ranked restoration as their customers' most frequent expected outcome category (Figure 2). For the second most frequent expected outcome category, 3 Support Engineers ranked restoration, 2 ranked troubleshooting, and 1 ranked root cause.

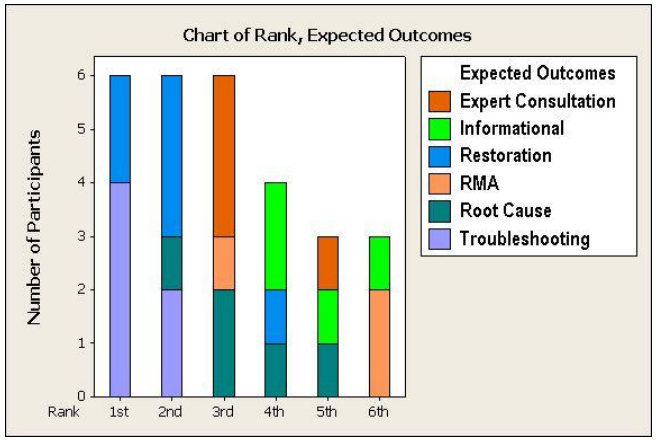

Figure 2: Support Engineers' Ranking of Outcome Categories

There were no additional categories or different category names identified by the 6 Support Engineers. We decided that the next study of the series should investigate whether customers can quickly identify one of the outcome categories for a service request.

Second Case Study — Can customers quickly select one category to match their expectations?

The second study evaluated how correctly the categories were worded to represent customers' expectations for their service requests.

Study Design

To evaluate the correctness of the category wordings, we used three questions:

- How quickly can customers select one of the outcome categories for a service request?

- How confident do customers feel about their selection?
- Do customers understand the predefined categories in relation to their expectations for a service request?

Two different evaluation methods were considered: Testing the correctness of the predefined category wordings as a baseline measure, or comparing the predefined category wordings with the words collected from the customers in the previous study. The latter could test the value of changing the design based on actual user linguistic behavior and was more in line with the low-fidelity approach of "Rapid Experimentation." We conducted a comparative usability study with 8 customer participants to compare two groups of category names listed in a simple design prototype: the names in Group One (Figure 3) were the predefined categories in the previous study; the names in Group Two (Figure 4) were collected from the customers in the previous study.

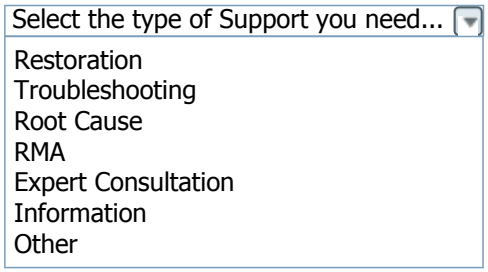

Figure 3: Category Group One

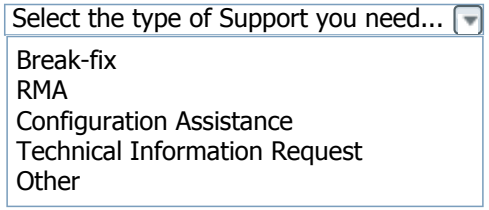

Figure 4: Category Group Two

We focused on collecting quantitative data, including the number of times participants selected the individual categories as their expected outcomes for their example service requests, the amount of time it took to make their selections, and the number of times they considered other categories before making the final selection. Some qualitative data were collected, including participants' comments on the accuracy of the categorization, the naming of the categories, and their definitions.

The participant criteria for this study were consistent with the customer criteria of the first study. We did not reuse participants from the first study.

Execution

The study took place in August, 2009. It required approximately 100 hours of work over 25 days. It took over 2 weeks to recruit 8 customer participants from 3 regions: 5 from the US and Canada, 1 from Europe, and 2 from Asia Pacific. The research team consisted of three part-time resources.

Each customer participant supplied three to five service requests that they had submitted to Cisco. A total of 33 service requests were used as examples for the study. During a 30-minute individual session, each participant first summarized what he/she expected Cisco to help with for a service request, and then selected a category from a group to match their expected outcome.

We split the 8 participants into two groups of 4. One group first used Category Group One (Figure 3) to make their requests and then used Group Two (Figure 4) to make the same requests. The second participant group did the same tasks but starting with Group Two followed by Group One. For each test we recorded the time it took for the participants to select a category and their confidence ratings.

During the participants' selection process, we recorded how many other categories the participants had considered before making their final selection.

After the category selection and ratings for the example service requests, the participants provided feedback about the two sets of categories, including how well the definitions resonated with their own understanding.

Results

The study disproved that customers could quickly select one category to match their expectations. We found that all participants had difficulty selecting a single category to match their expected outcomes for a service request. Neither version of the categories was easy for the participants to choose a single category to match their expectations.

Among the 33 service requests participants supplied (selecting a category 66 times from the 2 groups), there were 23 times when participants considered another category before making their final selection, which led us to believe that more than one category could represent their expected outcomes (Figure 5). We found that the participants regarded the RMA, troubleshooting, root cause, and restoration categories closely related, and that it was difficult for them to select just one category. Expert consultation and information were also hard to differentiate.

We learned that Group One was slightly easier for the participants to choose from, shown by higher confidence and less time for the selection (Figure 6). The category names and definitions needed to be improved to provide better distinction and cover more service request scenarios, such as licensing issues and interoperability questions. One category name in Group Two, "break-fix," was particularly difficult for participants outside of the US. We also learned that showing the definitions along with the category names was potentially helpful.

Group	Categories	# of Cases	# of Times Another Category was Considered
#1	Restoration	0	-
	RMA	4	1
	Root Cause	13	3
	Troubleshooting	9	4
	Expert Consultation	4	3
	Information	3	2
	Total	**33**	**13**
#2	Break-fix	13	4
	RMA	4	2
	Configuration Assistance	6	2
	Technical Assistance Request	6	2
	Other	4	-
	Total	**33**	**10**

Figure 5: Number of Times Another Category was Considered

Group	Categories	# of Cases	Average Confidence	Average Time (Sec.)
#1	Restoration	0	-	-
	RMA	0	-	-
	Root Cause	5	4.6	2.8
	Troubleshooting	4	5	1.7
	Expert Consultation	4	4	3.75
	Information	1	4	-
	Total	**14**	**4.4**	**2.75**
#2	Break Fix	6	3.5	7.25
	RMA	4	5	2.7
	Configuration Assistance	4	3.25	5.75
	Technical Assistance Request	2	3	9.7
	Other	3	5	-
	Total	**19**	**4.0**	**6.35**

Figure 6: Customers' Selection Performance for each Category

Based on the outcome of this study, the End-to-End Team added two design alternatives: allowing multiple selections of outcome categories, each with a "weight" factor (Figure 7); and allowing multiple selections of a primary outcome category and a number of secondary outcome categories (Figure 8).

We decided that the next study should investigate whether it would be easier for customers to select a primary outcome category and a number of secondary outcome categories.

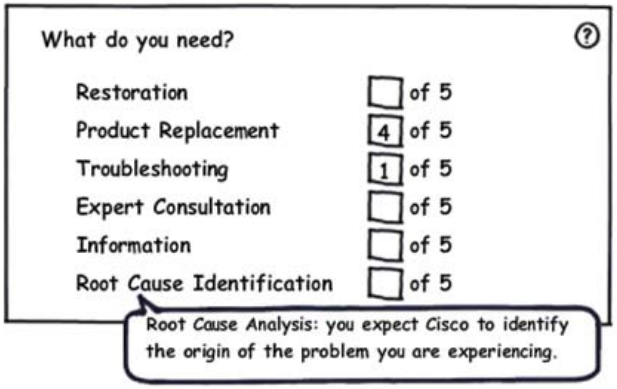

Figure 7: Allowing Multiple Selections of Outcome Categories each with a "Weight" Factor

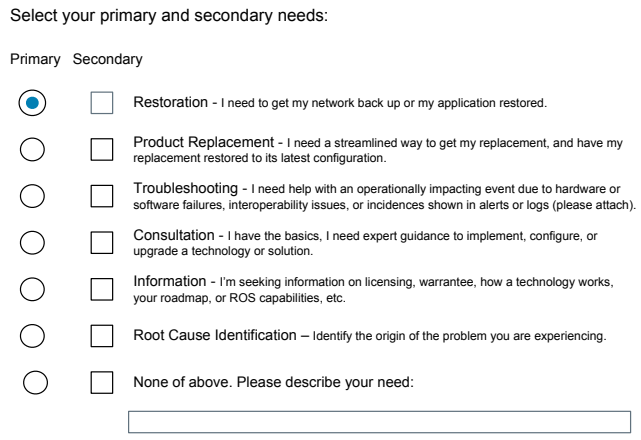

Figure 8: Allowing Multiple Selections of a Primary and Some Secondary Outcome Categories

Third Case Study — Can customers quickly select a primary outcome category and a number of secondary outcome categories to match their expectations?

After learning that customers had difficulty selecting a single category to match their expected outcomes for a service request, this study evaluated whether it would be easier for customers to select multiple outcome categories instead of a single one.

Study Design

This study explored whether a multiple-choice approach would be easier for customers, as well as for Cisco Support Engineers responsible for assisting customers in creating their service requests. We developed the following study questions:

- Overall, how easily did the participants select a primary outcome category and a number of secondary outcome categories for a given case?
- For each problem type, how easily did the participants select a primary outcome category and a number of secondary outcome categories? Is there a pattern in the outcome category selection?
- Among the outcome categories, how easily did the participants select a category based on the definition? How could the category names be changed to provide better clarity?
- How helpful or confusing were the definitions? How closely did they match participants' understanding of the outcome category names? How could they be changed to provide better clarity?

- Compared to the previous study, did the participants perform better or worse given the new multiple-choice concept? What were the main reasons for participants to select multiple outcome categories? Was it because they were confused by the categories, or were there indeed several expected outcomes?

The category names, definitions, and the selection method were designed to be easy for customers who open their service requests online, as well as for Cisco's frontline Support Engineers who use an online tool to assist customers who call Cisco to initiate their service requests. Results from the previous two studies suggested that we might need to create different category names and definitions for different regions. To eliminate this variable from the third study, we invited only customers and frontline Support Engineers from the US. We did not reuse participants from the previous two studies.

We conducted a usability study with 9 participants from the US: 5 customer participants and 4 frontline Support Engineers. The study collected experience data on the multiple-choice approach: the number of times participants selected the individual categories as their primary or secondary expected outcomes for their example service requests. It also collected participants' post-test ratings of the categories (ease of use and confidence in their selections), as well as helpfulness and clarity of the definitions. Some qualitative data were also collected, including participants' comments on the multiple-choice approach and on the wording of the definitions.

Execution

The study took place in September, 2009. It required approximately 75 hours of work spread over 19 days. The research team consisted of three part-time resources.

Both the customer participants and the frontline Support Engineers each supplied three to six service requests that they created in the past. A total of 33 service requests were used as examples for the study.

We asked each participant to email us their service request examples before their sessions. From the second study, we learned that if we had asked for the exact service request numbers from each participant, we could have further analyzed the service requests in relationship to the outcome categories selected. For example, each service request was assigned a problem code during the resolution process. We could analyze the relationship between the problem code assignment and the categories chosen for these cases.

During the 30-minute session, the participant first summarized what they expected Cisco to help them with in a service request. Then they followed the instructions in the design prototype and selected one or more categories to match their expected outcomes (Figure 8).

After the participant completed the selection for each service request, we asked "How confident are you that you have made the right selection(s) for this case?" and "How easy or difficult was it making your selections?" For each category selected, we then asked "How helpful or unhelpful did you feel the definition for [*each category selected*] was when you made the selection?" and "How clear or confusing did you feel the

definition for [*each category selected*] was when you made the selection?" We also asked participants for their reasons for the ratings, and if there were any suggestions for improvements.

After the category selection and ratings for the service requests, we collected feedback on the multiple-choice approach by asking "Why do you think you needed to select more than one area?" and "What if we only allowed you to select one choice?"

Results

The study confirmed that it was easier for customers and Cisco frontline Support Engineers to select a primary outcome category and a number of secondary outcome categories to match the customer's expectations for a service request. Participants had higher confidence using the multiple-choice design concept than the single-choice concept, and believed that the selection was fairly easy (Figure 9). All participants found at least one category to match their expected outcomes; none selected "None of the above."

Measure	Single Choice (Group One)	Single Choice (Group Two)	Multiple Choices
# of Participants	4	4	9
# of Cases	14	19	33
Average Confidence Rating (5 as Highest)	4.4	4.0	4.8
Average Ease-of-Use Rating (5 as Highest)	(Did not capture)	(Did not capture)	4.5

Figure 9: Customers' Ratings of the Selection Process for the Three Approaches

Participants selected multiple (from two to five) outcome categories for 24 of the 33 service requests, and selected only one outcome category for 9 service requests (Figure 10). Root Cause was never selected as the only outcome category.

Category	# of Times as the Only Category	# of Times as Primary of Multiple Categories
Restoration	1	3
Product Replacement	1	3
Troubleshooting	3	9
Consultation	3	4
Information	1	2
Root Cause	0	3
Total	**9**	**24**

Figure 10: Categories Selected as the Only Outcome Category or as the Primary of Multiple Outcome Categories

Despite the better ratings of the outcome category selection process, 3 of 5 customers commented that they preferred a single-choice selection because it would be easier and simpler. If the categories are distinctive and clearly defined, they would prefer to pick just one category. "I want a concise simple choice. Don't overload me with options." The other 2 customers, and all 4 frontline Support Engineers, preferred the multiple-choice approach, because there are different combinations of needs depending on the situations.

We learned that "Restoration" served as an urgency modifier to the customers, more than as an outcome category; "Information" was a simple type of "Consultation." When there is an issue, the customer and a Support Engineer troubleshoot the problem, during which they may identify the cause of the problem, then make the appropriate changes to restore functionality and prevent the issue from happening in the future. One design idea is to consolidate the six categories into three:

- Fix an issue with me — including "Troubleshooting," "Restoration," and "Root Cause Identification"
- Answer my question — including "Consultation" and "Information"
- Replace my product

A follow-on study (after the initial case study submission) has found that these three categories could cover a majority of customers' expectations for their service requests, and it was easier for them to select from these three categories than from six. The study also showed that it is necessary to provide subcategories to differentiate "Restoration" from "Troubleshooting" because "Restoration" indicates a higher urgency, and differentiate "Consultation" from "Information" because "Consultation" indicates a more complex question requiring a Support Engineer with more in-depth technical knowledge. These findings were consistent between English-speaking and non-English-speaking regions.

Lessons Learned

After the successful first case study, more small research teams were formed to explore additional conceptual experience improvement ideas. The "Rapid Experimentation" methodology applied equally well in other studies.

The more studies the research teams perform, the more skilled they become in using the methodology. The original goal was to complete two studies per month, with resources equivalent to four full-time persons. That goal was exceeded the first month. We now believe we can carry out as many as eight studies per month.

After completing more than 10 "Rapid Experimentation" studies, the authors have learned:

- Using the method can be challenging for individuals who are accustomed to larger-scale research projects seeking higher confidence levels. New team members require coaching and practice to stay in the rapid low-fidelity mindset.
- As the research teams work with internal subject-matter experts, they tend to lose sight of the original problem they are trying to solve and the study scope expands. The method requires strong management of scope and expectations.
- To narrow research questions and quickly feed the answers into the larger process, a well-defined "Rapid Experimentation" document template was very helpful: each study is to confirm or disprove one hypothesis using a simple "if…then…" statement; and all research questions should be tightly related to the hypothesis.

Conclusion

The iterative research findings described in this case study informed the End-to-End customer service experience team and designed a better mechanism for customers to select their expected outcomes, thus enabling Cisco to capture customer's expectations correctly up front and providing a personalized experience. Our goal is to eliminate at least 5% of the rerouting of all service requests, therefore reducing the average Time-To-Resolution. Not only will customer satisfaction increase, but so will the efficiency of Cisco's Support Engineers in fulfilling customers' requests, consequently saving on the Cost-Per-Call for Cisco.

Regular communication between the End-to-End Team and the research teams at every stage of each study—planning, design, early research data sharing, and reporting—enabled us to continuously inform or correct the End-to-End customer service experience design direction. Without these small and rapid studies, the End-to-End Team would have used more time to gather in-depth experience data for each design approach. Using resources roughly comparable to similar projects, we were 4 weeks ahead of the original 12-week schedule in achieving the same experience design goals.

The "Rapid Experimentation" method has proved to be a quick, effective, relatively low-cost approach for customer service experience research and design, with the nimbleness to allow for timely course corrections in the direction of a project.

Acknowledgements

The authors thank all members of the "Rapid Experimentation" research teams and the End-to-End Team of Integrated Service Delivery at Cisco who were part of the series of studies described here. We also thank Stephanie Rosenbaum and Laurie Kantner of Tec-Ed, David Curbow of Cisco Systems, and Dr. Brian Hoey of Roche, who reviewed and critiqued this case study before submission.

References and Citations

[1] Buchanan, Richard. Closing Keynote, Emergence 2007: Exploring the Boundaries of Service Design, Carnegie Mellon School of Design.

[2] Conboy K., Fitzgerald, B. Toward a conceptual framework of agile methods: a study of agility in different disciplines. Proceedings of the 2004 ACM workshop on Interdisciplinary software engineering research.

[3] Kroll, Per and Kruchten, Phillip. The Rational Unified Process Made Easy: A Practitioner's Guide to the RUP. Addison-Wesley Professional (2003).

[4] Medlock, M. C., Wixon, D., Terrano, M., Romero, R. L., Fulton, W. 2002. Using the RITE method to improve products; a definition and a case study. In Proceedings of Usability Professionals Association, Orlando, Florida, USA, 2002.

[5] Poppendieck, Mary and Poppendieck, Tom. Lean Software Development: An Agile Toolkit. Addison-Wesley Professional (2003).

[6] Spear, Steven J. and Bowen, H. Kent. Decoding the DNA of the Toyota Production System. Harvard Business Review (2009).

[7] Spear, Steven J. Chasing the Rabbit: How Market Leaders Outdistance the Competition and How Great Companies Can Catch Up and Win. McGraw-Hill (2008).