

Who Needs Release and DevOps Engineers, and Why?

Noureddine Kerzazi
ENSIAS, Mohammed V University in Rabat
n.kerzazi@um5s.net.ma

Bram Adams
MCIS, Polytechnique Montreal
bram.adams@polymtl.ca

ABSTRACT

The recent surge in interest in continuous delivery has opened up the job market for release and DevOps engineers. However, despite an increasing number of conferences and publications on continuous delivery, smaller companies and start-ups still have a hard time determining the core tasks their future release and DevOps engineers should be responsible for (and what the differences between those two roles are), while universities are not sure what essential techniques and skills they should teach to their students. This paper performs an empirical analysis of online job postings to determine and compare the main tasks of release and DevOps engineers, globally and across countries. Our qualitative analysis shows that automation is the most important activity across the three roles, as articulated in job posting description data, and that the release engineer role combines the top activities of the DevOps and more traditional build engineer roles. Finally, different countries have a moderate degree of similarity between their ads, although each country has its specific focus.

CCS Concepts

•Software and its engineering → Software configuration management and version control systems; Software version control;

Keywords

devops, release engineer, job description, empirical study

1. INTRODUCTION

Ask two people what a developer does and you will get pretty much the same answer: a developer writes code. However, only little consensus can be found in researcher and practitioner circles on what a release [3, 8] or DevOps [4] engineer is supposed to do. Do these roles refer to a development team member with some kind of overarching technical responsibility towards a software release? Is it a managing

role overseeing and planning the transfer of deliverables towards the production environment? Penners et al. [7] found that the experts consulted by them did not share a unique interpretation of the DevOps and release engineering roles, while Bass et al. [4] even define DevOps in terms of release engineering, i.e., “*DevOps is a set of practices intended to reduce the time between committing a change to a system and the change being placed into normal production, while ensuring high quality.*”

The Capability Maturity Model Integration (CMMI) describes Release Management as emphasizing the Product Integration (PI) process area [12]. The purpose of PI is to assemble the product from product components, to ensure that the integrated product works properly, then to deliver it. In a nutshell, PI practices are: plan the release, test and verify the release, and deliver the release. Within another constellation of “best” practices, the IT Infrastructure Library (ITIL) defines Release Management as a process including people, functions, systems and activities to plan, package, build, test and deploy software and hardware releases effectively in the production environment [9]. However, neither the CMMI nor ITIL governing bodies of best practices have provided clear guidelines about how to effectively manage a software release process and how to harmonize the release process definitions.

Instead, large software companies like (amongst others) Google, Facebook, Mozilla and Netflix pioneered and invested substantially in modern release engineering techniques and technology, enabling consistently short release cycle times of one month down to one day or even several hours. To accomplish this, as well as to optimize the post-release management of bugs and other quality issues, the traditional roles of build engineer, sysadmin and tool developer were merged and reinvented into new software roles like “release engineer” and “DevOps engineer”, conceived with a mission to streamline the handoffs from development to operations.

Given the impact that release and DevOps engineers have had on these large companies, smaller companies, with less resources, as well as start-ups are now trying to jump on the bandwagon. However, they do not have the luxury of being able to experiment with these roles, furthermore they need to piece together the expected activities and responsibilities of said engineers from what literature and online blog posts and videos can tell them. Schools and universities are in a similar situation, since they want to prepare their students for these new software engineering roles, but need to improvise regarding the curriculum they should be offering.

This paper analyzes online job postings to (1) define the

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CSED '16 May 14–15, 2016, Austin, TX, USA

© 2016 ACM. ISBN 978-1-4503-2138-9.

DOI: 10.1145/1235

scope of the release and DevOps engineering activities used by current adopters, and (2) compare these activities, both globally and between different countries. We find that:

- the most important activity across the three roles is scripting of otherwise manual steps;
- the role of release engineer seems to combine the most important activities of DevOps and more traditional build engineers, either because of incorrect choice of role name or roles taking up more responsibilities than would be expected from them;
- different countries can have a different focus, even though there is a moderate agreement between job ads across countries.

The remainder of this paper is structured as follows. section 3 introduces the design of our empirical study, based on online job advertisements, including the research questions, and followed methodology. Then, the results are discussed in section 4. section 5 provides a broader discussion of our findings, while section 6 presents their limitations. After surveying related work in section 2, section 7 concludes the paper and sketches future work.

2. RELATED WORK

The closest work to this paper is the work by Kanij et al. [10] on collecting job descriptions of software testers from a job ad website, over a period of five days, aiming to distill a list of software testing responsibilities. The authors found that the test-specific responsibilities are divided into several unit tasks, including test suite generation and execution of test plans. Apart from test-specific responsibilities, software testers must also perform many other duties such as debugging, planning, maintenance, management, and collaboration with others roles.

Penners et al. [7] worked on defining the roles of DevOps and release engineers. We share their motivation that the responsibilities and differences between both roles currently are unclear, which has a negative impact on companies trying to understand these roles, as well as students trying to prepare for them. However, the main difference between this work and ours is the approach followed. While we empirically study online job ads (and also consider a third role of build engineer), Penners et al. contacted experts in the field and tried to distill definitions for the roles from the experts' replies. Penners et al. conclude that:

Release engineering is a software engineering discipline concerned with the development, implementation, and improvement of processes to deploy high-quality software reliably and predictably.

This definition confirms our main findings regarding release engineering, although we found that the improvement aspect (pipeline optimization activity) in fact was more popular in DevOps job ads. Furthermore, the aspects “Reliably” and “Predictably” are hard to measure textually from job ads, hence we did not really find indications of these. DevOps is defined by Penners et al. as:

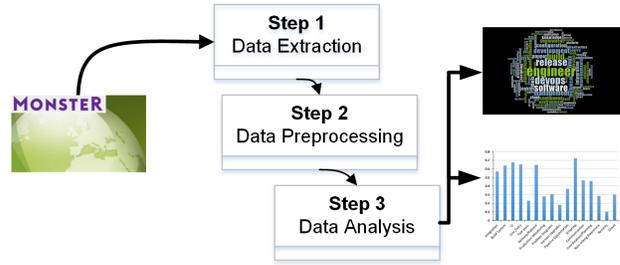


Figure 1: Overview of the approach.

an organizational approach that stresses empathy and cross-functional collaboration within and between teams – especially development and IT operations – in software development organizations, in order to operate resilient systems and accelerate delivery of changes.

The aspect of “collaboration” indeed appears as a top activity in DevOps ads, but also for release engineering ads. As mentioned, pipeline optimization indeed turned out to be a major activity for DevOps ads. Overall, we found correspondences with the definitions of Penners et al. (and Bass et al. [4]), but also some unexpected differences.

A roundtable with three release engineers [2] discussed the actual challenges of release engineers. For instance, a release engineer from Google stated that release engineering is often not even mentioned in courses where it should be mentioned, mainly because the “*release engineering practice itself has been hard to define. As you see from the answers of your other guests, the approaches are quite diverse in nature and scope.*” This motivates our research for further comparison between the different release engineering roles, including regions, scales, and nature of organizations.

3. RESEARCH SETTING

This section presents our methodology for studying what organizations are looking for when posting a job ad for release or DevOps engineers. In particular, our study addresses the following questions:

1. What are the Core Activities of the Two Roles?
2. How do Core Activities Differ across Geographical Regions?
3. Who Needs Release and DevOps Engineers, and Why?

Our research method is summarized in Figure 1. It consists of three steps: data extraction, data preprocessing and data analysis. This section describes each of these steps in more detail.

3.1 Step 1: Data Extraction

Job ads contain textual descriptions of expected responsibilities, qualifications and bonus skills required from a prospective release or DevOps engineer. For example, a job ad could contain text like “. . . *The Release Engineer will be responsible for managing complex code builds and supervising day-to-day code integration activities for a team of developers spanning multiple applications, environments and locations. The*

candidate will be responsible for developing and maintaining portions of the integration and release processes within our team. The Release Engineer will be responsible for creating and maintaining environment configuration and controls, code integrity, code conflict resolution, and will work closely with software development, programming and QA teams to review code for compatibility issues, resolve issues as they arise and implement deployment processes . . . ”.

In September 2015, we collected one month worth of job ads using the search query “DevOps OR Release Engineer” from monster.com [11], focusing on five countries. This site is one of the largest online employment platforms in the world. Given that we found several thousands for the US alone, we decided to sample the ads to make qualitative (i.e., manual) analysis feasible. In order to make meaningful conclusions, our set of sampled ads needs to be “sufficiently large”. To achieve a confidence level of 95% and confidence interval of 10% for our sample set¹, we randomly selected 211 ads, spread across five top software engineering countries, i.e., USA (119 items), Canada (33), UK (49), Australia (6) and India (4). We stored the metadata of the collected job ads in csv files, while the job ad descriptions themselves were downloaded as text files.

3.2 Step 2: Data Preprocessing

After gathering the data, we extracted the title, company and region requesting the job, the date of posting, expected salary, and a reference for a detailed description (i.e., post resume) containing responsibilities, expected skills and duties, degree among other, in an unstructured format.

The job extraction required three additional steps:

1. **Filtering** - A job posting is an advertisement meant to attract job applicants. Hence, it tries to make both the job and the organization stand out by appealing to the interests and preferences of the specific people organizations want to attract. This means that we had to filter out all irrelevant information within the description of the job.
2. **Classification** - We then classified job posts according to their title into the release or DevOps engineering role. Although our primary search query did not include the role of build engineer (traditional role of build process responsible), it retrieved many jobs with a job title such as “Build Release Engineer”, since this role is closely related to those of release and DevOps engineer. Instead of ignoring build engineer job ads, our study included them to help us put our results for DevOps and release engineer into context.
3. **Lemmatization** - We used a checklist to cluster the many technical terms and technologies that have the same purpose. This also includes merging different inflected forms of words that are syntactically different but functionally equal. This step reduces considerably the effort of manual inspection later. For instance, the terms “Jenkins”, “Buildbot”, and “MS-Build” are grouped into the activity “Continuous Integration”. Another example would cluster “version control, Git, SVN, Mercurial, ...”.

¹This means that if we would find that 20% of the sampled job ads are related to a particular activity, we should interpret this as 20±10% of the whole population of job ads.

3.3 Step 3: Data Analysis

We start our data exploration by simple textual analysis such as word clouds. Figure 2 shows the resulting word cloud across the 211 job ads, with the most common terms across the ads shown in larger typeface than less popular terms. Looking at the word cloud, we see that several dominant concepts emerge, such as Release, Build, DevOps, and Scripting. We also built more advanced topic models using Latent Dirichlet Allocation (LDA) to cluster terms into semantically more meaningful topics representing concrete activities of build/release/DevOps engineers.

We then combined the dominant activities found by the topic model with the activities and responsibilities of build/release/DevOps engineers known from literature [4, 8] and from specialized events or conferences on these topics [1, 3]. In particular, we performed a Task Analysis, which is a methodology used to study the activities that an operator needs to perform in order to accomplish a goal [5]. Put briefly, it involves the detailed analysis of manual, automated, and cognitive activities, task durations and frequency along with environmental conditions. The output of the Task Analysis is a list of core activities required to complete a software release adequately. We then structured these core activities in an initial checklist to perform the analysis and classification of job ads. During the qualitative analysis of the job ads (see below), the resulting activities and responsibilities were refined into the 16 activities described in Table 1. These activities can be divided into core activities (white background) and supporting activities (grey background).

To understand the differences between the activities of build/DevOps/release engineers, we manually analyze 144 out of the 211 job ads using qualitative analysis to gain further insight into the nature of responsibilities versus duties as expressed by practitioners. Basically, the two authors manually read half of these 144 job ads, tagging sentences by one or more of the activities of Table 1 that were predetermined from the topic model, literature or events. In case a new kind of activity arose, it was added to our checklist, eventually giving rise to Table 1. The outcome was a spreadsheet with one row per job ad and an “X” for each applicable activity.

Job ads were mapped to a role by searching the ad titles for the terms “Build”, “DevOps” and “Release”. Ads were also mapped to a country based on the corresponding field on the monster.com website. As such, we could group each ad per role and/or per country. Table 2 shows the distribution of job ads across the three roles.

4. RESULTS

This section discusses the results of our analyses.

4.1 Build vs. DevOps vs. Release Engineers

Figure 3 shows that scripting (automation) is the most important activity across the three roles. 72.2% of the analyzed job ads require automation of release engineering steps through scripting. This activity is closely followed in popularity by continuous integration (67.4%), environment/infrastructure (65.3%), delivery/release (64.6%), build system (63.9%) and integration (56.9%). Especially the environment/infrastructure scores are surprisingly high, given that programming languages for infrastructure-as-code

Table 1: Core/Supporting Activities Identified in the job ads.

Themes	Description
Integration	Source control management (SCM), including branching and merging strategies for parallel development as well as using SCM tools like git and subversion.
Build System	Techniques and tools for building and packaging source code and other files into deliverable (e.g., Ant, Maven, Makefile).
Continuous Integration	Managing automated build and quality assurance tools on certain events using tools like Jenkins and MsBuild. For example, after each change committed by a developer, the resulting deliverable should be built and tested on a dedicated server. This also includes the tagging and storing of build artifacts in artifact repositories like JFrog or Nexus.
Environment/Infra	Defining and managing the environments and infrastructure (servers, VMs or containers) for different staging activities, such as development, testing, pre-production and production (dev, test, pre prod, and prod). This also involves coordinating with other roles.
Test Execution	Building the test harness and/or scaling up test execution in larger build systems.
Delivery/Release	Setting up and maintaining the pipeline for deployment and release of a new version of a software project.
Operation Monitoring	Monitoring and surveillance of the production environment.
Problem Diagnosis	Diagnosis of post-release defects and systems behaviour.
Version Upgrades	Upgrading an existing production environment to a newer one, possibly using techniques like canary deployment or rollback.
Pipeline Optimization	Optimizing the different activities of the release process, in particular to move towards continuous delivery.
Scripting	Scripting and automation of manual release engineering tasks, typically using scripting languages like Bash, Python or PowerShell.
Communication	Ability of conveying one’s message to other roles in the form of informal discussions, official presentation, documentation, email or any other communication medium.
Coordination/Planning	Supporting the planning of a release such as determining the release’s roadmap and cycle time. This requires explicit coordination with other team members, such as developers, database administrators and IT staff.
Non-releng Awareness	Involvement in other development cycle activities, such as developing, architecting or testing.
Security	Need for making release infrastructure and processes security-aware.
Cloud	Knowledge of and involvement in cloud-based build, test or deployment environments.

such as Chef and Puppet are still relatively young. The other top activities are more traditional release activities.

On the other hand, the security (10.4%), software upgrade (18.1%) and test execution (22.9%) activities were the least popular across the analyzed job ads. Since security is a relatively recent activity in the context of release engineering or DevOps [4], its score is actually quite high. This does not hold for the software upgrade activity, which is a mainstay of the job of sysadmin (one of the roles out of which DevOps has grown). Either the companies of the analyzed ads already had tools or employees in place for performing upgrades, or they forgot to explicitly mention this activity. Finally, the low score of test execution could be explained by its main focus on speeding up test runs, which might be subsumed by the pipeline optimization activity, which is moderately popular (36.8%).

Figure 4 shows that release engineer job ads share the top activities of build and DevOps engineers. If we consider the top 5 activities of the three roles, we find that build system, scripting and delivery/release are top activities in all three roles. Whereas build engineers add CI (92% of build engineer ads) and integration (76%), DevOps engineers add environment/infrastructure (72.6%) and CI (64.4%). Release engineer job ads incorporate all of these, i.e., environment/infrastructure (63%), integration (60.9%) and, slightly outside the top 5, CI (58.5%). The

activities of build/DevOps engineers correspond to what one would expect, in particular the DevOps engineer’s focus on infrastructure-as-code.

Pipeline optimization (49.3%) and cloud (42.5%) are two important activities for DevOps engineers only, while problem diagnosis (41.3%) is a major activity for release engineers, and production monitoring is important for both roles. Other activities are more similar, in particular scripting, communication (up to 54.3% for release engineers) and most of the other top 5 activities. Surprisingly, problem diagnosis is less important to DevOps job ads than to release engineering job ads. Conversely, optimizing the release pipeline to enable more rapid releases was more popular in DevOps ads than in release engineering ads. We suspect that these observations might be due to the emphasis of the DevOps role on post-release operations rather than on software releases. Furthermore, release engineers effectively could be playing the role of DevOps engineer in practice, as was suggested earlier from our analysis of top 5 activities.

4.2 Geographical Differences between the 3 Roles

US DevOps engineers are expected much more frequently to perform automation (100%) and coordination/planning (86.7% as shown in Figure 6) than their UK and Canadian counterparts. Note that au-

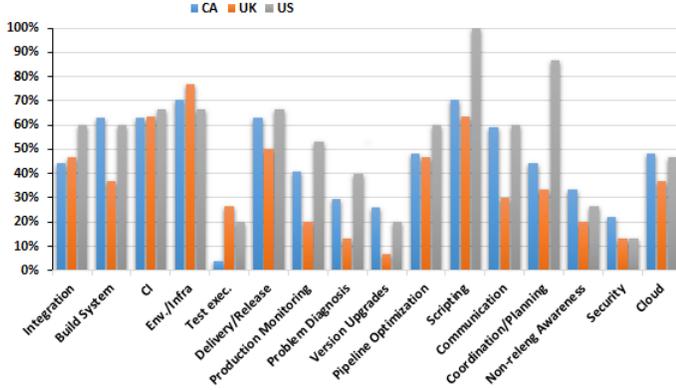


Figure 5: Popularity of the 16 activities per country for the DevOps role. Percentages are relative to the number of DevOps ads of a given country.

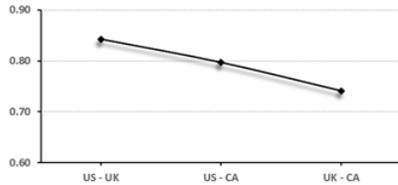


Figure 6: Pearson correlation coefficient of the word similarity between the job ads of each pair of countries.

5. DISCUSSION

5.1 Who Needs Release and DevOps Engineers, and Why?

Although not an explicit variable in our analysis, we observed that build/DevOps/release engineers are needed by any kind of company, regardless of whether it is a start-up or an established one (based on the job ad description). For start-ups, the responsibilities are more guided towards supporting development, followed by build, continuous integration, and fast delivery (building an effective pipeline of releases). This is shown in Figure 7 as a red circle, collecting all job ads mentioning development-related terms.

Expanding companies focus more on configuration management, likely to support parallel development, on explicit coordination (Bamboo, Jira, TFS) and, most importantly, management of environments for development, testing and production (blue circle). Finally, established companies require production monitoring, troubleshooting, and general improvement of release practices (green circle). These findings are based on our observations during manual analysis, but more explicit analysis is required to fully quantify the link between the status of a company and the 16 activities of Table 1.

5.2 The Three Roles are Crosscutting

Regardless of the role, our findings show that build/DevOps/release engineers need to perform a wide variety of tasks spread across the typical release pipeline, i.e., their role is crosscutting in nature. This is also clear by the importance of coordination/planning and communication. Indeed, while

Table 2: Distribution of job ads across the three roles.

Region	DevOps	Release Eng.	Build Eng.	Total
US	41	50	28	119
CA	25	2	6	33
UK	30	8	11	49
AU	5	1	0	6
IN	1	0	3	4
Total	102	61	48	211

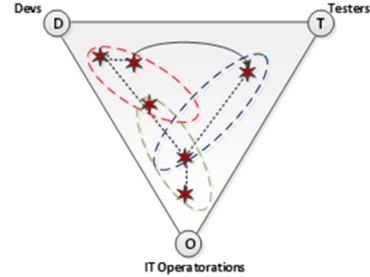


Figure 7: Ternary representation of release management activities.

developers and testers often are limited to specific architectural components of a system, the three roles are responsible for combining the components into a consistent whole, deploying and operating it. As expected, this requires a thorough insight into the overall architecture [4], where the three roles often play the role of arbiter to resolve integration and team conflicts. In contrast to this high-level responsibility, the three roles also need to perform very technical tasks like automation and scripting.

While it may sound that release engineers focus primarily on operational views, this is not the case. There are many other roles that form part of the release team, some of whom have not traditionally considered themselves to be part of it. For instance, testers play a large role in developing effective automated test infrastructure. Involvement of product owners ensures that capacity needs as well as anticipated costs and continuity of services will be considered early on in the design of scaling. The growing observed demand on DevOps has proven that this relatively new role helps guarantee that release practices will serve the purposes of the organization.

6. THREATS TO VALIDITY

As for each empirical study, there are certain choices in the data extraction and preparation that might impact the findings (construct validity of the study). One major assumption is that job ads are accurate, since our manual analysis is based on their textual description and we also automatically queried them for terms like “Release engineer” or “DevOps”. Since job ads are meant to find the right person for a job, we believe that, barring collective misunderstanding of the actual responsibilities of one of the three analyzed roles, the mixture of job ads from different companies and countries overall should filter out most of the noise in the data.

Manual analysis helped us to find the right interpretation of certain activities. However, the qualitative evaluation of the topic relevance to the job duties was done by the authors of the paper and not by recruitment experts (i.e., human resources). This is a threat to validity as the evaluators could

be biased or could have incomplete knowledge about specific activities that are required with respect to the context, platforms, and scales. Another threat to validity relates to the level of disagreement between coders of what constitutes an adequate set of topics to be considered for classifying activities. We tried to alleviate disagreement by deriving and using a checklist of core activities based on literature.

Furthermore, we only mined the monster.com website. Although this is one of the major job ad web sites worldwide, other websites exist, some of which might be country-specific and provide more accurate reflections of the job ads of a particular country. More empirical studies on other ad sites are necessary to mitigate this threat.

Furthermore, alternative explanations for our findings could be found because of variables not considered (internal validity of the study). We briefly discussed the role of the status of a company, but the role played by the company's country (e.g., mostly outsourcing) could also impact our findings. For example, release engineers might be hired in one country and sent to work in another one.

Finally, our findings might not be generalizable to other companies, job ad web sites or countries (external validity of the study). As mentioned earlier, additional studies are necessary, both on job ad data, as well as other data sources (for example as done by Penners et al. [7]).

7. CONCLUSION

The lack of a common vocabulary and a body of knowledge on DevOps and release engineering negatively affects the implementation and training of these roles for companies and universities. Companies that do not fully understand these roles, do not know exactly what skills to look for when hiring people to fill these relatively new positions. At the same time, universities struggle to design a suitable curriculum for these roles, mostly trying to shoehorn these topics into existing courses.

This article performs an empirical study on online job ads to try and learn from existing companies' job requirements for these roles. Through a mixture of approaches, we identify 16 core activities that we manually looked for in the sampled job ads. We found that scripting (or, more generally, automation) is the most important activity across the three roles, and that release engineering job ads seem to combine the top activities of build and DevOps engineers. This might be because of incorrect choice of role name, or release engineers taking on more general responsibilities than strictly release engineering-related ones.

Finally, even though there is a moderate agreement between job ads across countries, we found some country-specific differences. For example, automation and coordination/planning are stressed more in US job ads than in UK or Canadian ads. More empirical studies are necessary to fully understand the differences between the three roles.

In particular, we are currently performing a survey amongst release, DevOps and build engineers to complement our current findings.

8. REFERENCES

- [1] B. Adams, S. Bellomo, C. Bird, F. Khomh, and K. Moir. 2nd international workshop on release engineering (releng' 14). In *(Releng' 14)*, at Google Mountain View, CA, USA, 2014.
- [2] B. Adams, S. Bellomo, C. Bird, T. Marshall-Keim, F. Khomh, and K. Moir. The practice and future of release engineering: A roundtable with three release engineers. *Software, IEEE*, 32(2):42–49, Mar 2015.
- [3] B. Adams and S. McIntosh. Modern release engineering in a nutshell – why researchers should care. In *Leaders of Tomorrow: Future of Software Engineering, Proceedings of the 23rd IEEE International Conference on Software Analysis, Evolution, and Reengineering (SANER)*, Osaka, Japan, March 2016.
- [4] L. Bass, I. Weber, and L. Zhu. *DevOps: A Software Architect's Perspective*. Addison-Wesley Professional, 1st edition, 2015.
- [5] A. Crystal and B. Ellington. Task analysis and human-computer interaction: approaches, techniques, and levels of analysis. In *Proceedings of the Tenth Americas Conference on Information Systems*, New York, Aug. 2004.
- [6] A. Dyck, R. Penners, and H. Lichter. Towards definitions for release engineering and devops. In *Proceedings of the Third International Workshop on Release Engineering, RELENG '15*, pages 3–3, Piscataway, NJ, USA, 2015.
- [7] J. Humble and D. Farley. *Continuous Delivery: Reliable Software Releases Through Build, Test, and Deployment Automation*. Addison-Wesley Professional, 1st edition, 2010.
- [8] ITIL. Information technology infrastructure library (edition 2011). Technical report, Office of Government Commerce ITIL, 2011.
- [9] T. Kanij, R. Merkel, and J. Grundy. An empirical study to review and revise job responsibilities of software testers. In *IEEE Symposium on Visual Languages and Human-Centric Computing*, pages 89–92. IEEE CPS, 2014.
- [10] Monster.com. Online recruitment website. <http://www.monster.com>, 2016. Accessed: 2016-01-02.
- [11] SEI. *CMMI for Development, Version 1.3*. Software Engineering Institute, www.sei.cmu.edu/reports/10tr033.pdf, CMU/SEI-2010-TR-033. pp-469, 2010.