



Introduction

- Companies depend on online availability, reducing downtime is essential
- Learning from previous incidents to more efficiently solve future incidents
- Root cause analysis on 258 incident reports from Google
- Using AI to extract the relevant details
- Linking key numbers and extracted information to TTR to reveal patterns

Related Literature

- Many incidents caused by small quickly noticeable changes
- Number of available skilled engineers is a major factor in determining the MTTR
- Proven categories of root causes used in this study

Research Question

To what extent do different types of software changes correlate with the mean time to repair (MTTR) of incidents at Google?

Methodology

1. Downloaded 979 incident reports from Google's public repository
2. Filtered to reports containing root cause information (n=258)
3. Preprocessed text using NLP techniques and data cleaning
4. Trained classification models for automated change type identification
5. Statistical evaluation of change impact on MTTR and TTD metrics

Results

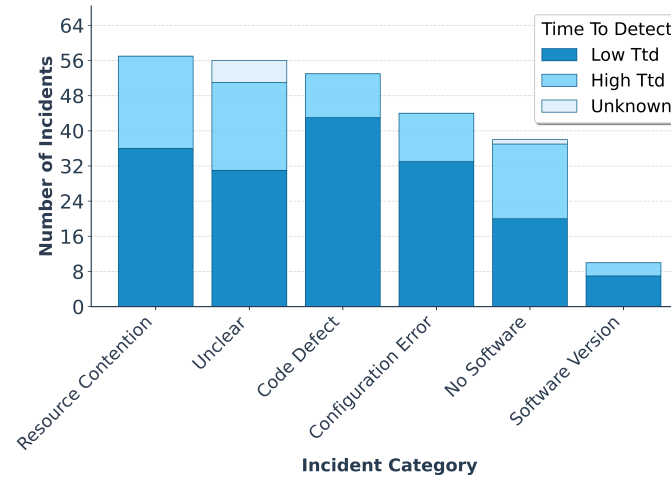


Figure 1: TTD distribution per root cause category

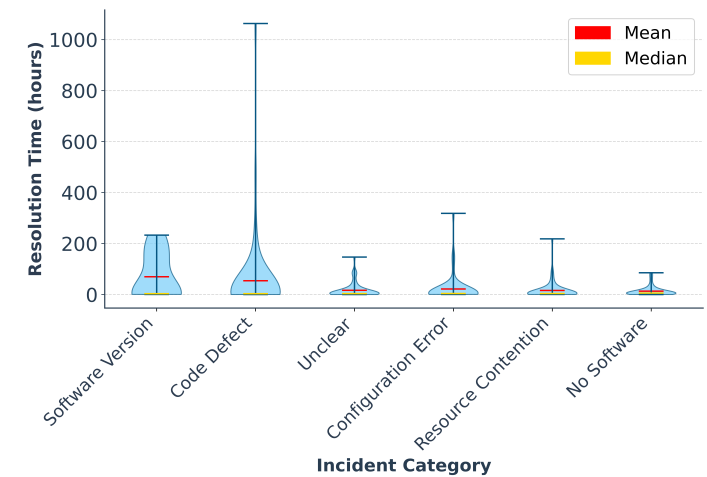


Figure 2: TTR distribution per root cause category

Key Findings

- **Software version** changes (69.8 hours) and **Code defects** (54.0 hours) have significantly higher MTTRs than average 26.7 hours (Figure 2)
- **Code defects** show the highest standard deviation for resolution time of 163.58 hours (Figure 2)
- **Code defects** and **Configuration errors** have at least 75% of their incidents categorized as low TTD (Figure 1)
- **Software version** incidents are the least common incidents, but do exhibit the highest MTTR (Figures 1 & 2)

Discussion

Limitations

- Limited dataset from only one company (Google)
- Potential selection bias in publicly available reports

Key Conclusions

- **Software version incompatibilities** represent the highest-risk change type with 69.8h average MTTR
- **Code defects** and **configuration errors** are quickly detected but varied in resolution time

Future Work

- Validate findings across multiple companies and industries
- Expand analysis to include incident severity classifications

- Use **insights** from past incident to improve system reliability
- Organizations should **tailor response strategies** based on root cause patterns