



eBook

# How Debugging Is Changing



# Table of Contents

<b>3</b>	<b>Introduction</b>
<b>4</b>	<b>Part 1: Testing Alone Isn't Working</b>
<b>5</b>	<b>Part 2: Imperfect Debugging</b>
<b>6</b>	<b>Part 3: Strategies to Improve Debugging</b>
6	Automation
6	Log Aggregation
7	Observability
<b>8</b>	<b>Part 4: Debug Faster with Continuous Code Improvement</b>
8	Grouping
8	Contextual Information
9	Connecting Errors Across All Environments
<b>10</b>	<b>Part 5: Conclusion</b>

# Introduction

[Up to 50% of developer's time](#) is spent finding, triaging, and fixing bugs rather than working on new features or deploying updates. Since a development team's business value is often measured by how frequently and how fast new features are successfully deployed to production, the amount of time spent on bugs can be troubling.

And as software becomes more complex and the systems that deploy it become more diverse, the goals of better software quality and less time spent debugging become even more challenging. For most organizations, errors and debugging are a significant drag on the software development lifecycle (SDLC).

We are going to take a look at the problems in debugging software errors as a way to understand how to minimize the costs and time of debugging. We'll look at why many software testing strategies fail, a few alternate strategies that can improve debugging processes, and specifically how tools like error monitoring and other components of continuous code improvement can provide an improved understanding of debugging issues. The goal is that by understanding debugging across all stages of the SDLC, there should be fewer bugs in your code and they should be easier to resolve, and fixed sooner, before they have an impact in production.

# 1 Testing Alone Isn't Working

Software testing is supposed to catch defects and provide the information needed to debug and fix them before they make it to production. Developer unit tests catch logic issues, continuous integration/continuous deployment (CI/CD) tools run integration tests, QA executes functional and automated tests of use cases, and users give software a final check during UAT. In practice, however, software is usually much more complicated than anticipated, and this process has many failure points.

Here are some examples of the complications:

- ✓ A developer's local environment rarely mirrors the complexity of production, and his or her unit tests can miss errors that become obvious in production.
- ✓ Testers rarely have the time or budget to create 100% coverage, leaving holes in use case testing.
- ✓ A reasonably-sized code base interacts with many other services, both in-house and external, injecting unknowns and dependencies into the software.
- ✓ Modern software stacks are much more complex than their monolithic counterparts from years ago. Microservices, cloud integration, and CI/CD components interact with the code in multiple ways making it extremely challenging to test.

As a result of this complexity, bugs often slip through the process and into production. Once those bugs are found, it's challenging to dig through the complexity, variation in environments, integrated tools, and cloud-based deployments to find the information to quickly resolve the bugs.

Debugging these issues costs time and, invariably, costs the business a lot of money. [A study by the University of Cambridge](#) noted that software debugging costs businesses \$312 billion per year and growing. With the advent of more complex software stacks, cloud-based infrastructure, and microservices integration, this number will grow as testing itself becomes more complex and difficult.

## 2 Imperfect Debugging

Although helpful in weeding out bugs, the traditional software testing process is part of what is known as [imperfect debugging](#). Imperfect debugging is not ideal, as it can:

- ✓ Increase the possibility of introducing other failures
- ✓ Fail to resolve the identified failure
- ✓ Decrease the mean time between software failures (MTBSF)

On the other hand, debugging is considered perfect when:

- ✓ Debugging corresponds to an improvement to the software's reliability
- ✓ The failure is corrected
- ✓ The MTBSF increases

The effect of imperfect debugging has a direct correlation on software cost and a negative effect on software release time. In this chart based on the IEEE Study "[A Study of the Effect of Imperfect Debugging on Software Cost](#)", we can see the relationship of testing level  $p$  and software cost  $C$ :

### Change of Software Cost with Testing-Level $p$

$p$	$C(\$10)$
0.50	65.257
0.55	57.132
0.60	51.023
0.65	46.268
0.70	42.455
0.75	39.315
0.80	36.674
0.85	34.410
0.90	32.442
0.95	30.708
1.00	29.166

## 3 Strategies to Improve Debugging

Now let's look at three different strategies that can be used to move towards more perfect debugging—automation, log aggregation, and observability.

### Automation

Automation practices that bridge the gap between developer and operation teams, such as CI/CD, can in some instances make it more difficult to track down issues. However, the same tool that causes these issues can also lead to earlier bug detection and mitigation. Tools such as [Jenkins](#) and [CircleCI](#) encourage the goal of consistent and automated builds and tests. This encourages teams to commit code more frequently, which leads to the concept of continuous testing—automated regression and other tests run as part of the CI/CD process.

Automated tests, using methods such as [Test-Driven Development \(TDD\)](#) or [Behavior-Driven Development \(BDD\)](#), and running more frequently on smaller sets of code changes, make it much easier to isolate changes that could have caused defects. Less code has changed, fewer people have touched the code, and fewer features were implemented. Therefore, the process of finding the root cause becomes faster and simpler.

To get even more usefulness out of these automation tools, they can be moved as close to the developer's unit-testing environment as possible. This is known as "shifting left." Shifting left espouses testing earlier in the SDLC on smaller amounts of code, leading to a lower mean-time-to-fix (MTTF) for potential

bugs. The more testing drifts in the wrong direction (to the right and closer to launch) the higher the likelihood that detected bugs will move to technical debt instead of being addressed.

### Log Aggregation

Log aggregation is another method used in improving the debugging process. The collection and analysis of log data provides operational visibility across the stack. Log aggregation, allows for recording and even possibly replay of events leading up to a software crash or error. Log management and features such as search capabilities, data visualization, and programmable event triggers lead to a holistic view of events, traces, and event information in code.

However, while log aggregation is useful, it can be costly and hard to standardize. With increased event tracking or additional events triggered by always-on systems such as microservices or IoT devices, businesses face the challenge of rising data storage costs, integration difficulties, and the need for more advanced querying technologies against the reising volumes of data.

But, even if most information does get logged, the equally challenging task of making sense of all that data arises. Aggregation, collection, and analysis of log data have been common practice, but it still isn't enough. Intelligence around this aggregation plus integration with your other business tools (such as ticketing systems) are necessary to realize its full business value. Tools that can automate and make intelligent decisions about aggregation and grouping can really simplify this task.

## Observability

Observability – the measure of how well the state of the stack can be inferred through knowledge of its outputs – is another strategy for debugging. Observability means that you have visibility into the system’s behavior to better understand what’s happening across the stack. If logging shows what the stack is doing, observability shows why it is doing it. In observability, manual querying and step debugging are replaced by correlation analysis and advanced analytics.

One newer concept in observability is storing every state of every debugging step into an indexed, searchable database. These tools store every exception, warning, error, and log entry in a data system typically using schema-less, search engine frameworks like Lucene in Elasticsearch. This is referred to as omniscient debugging and it holds promise of quickly identifying and historically tracing known patterns through a combination of machine learning and data science/engineering technologies.

However, according to a recent IDC study, 50% of all organizations process more than 100GB of data a day. An omniscient debugging system would increase this requirement many fold.

Coupled with the fact that in a world where more devices are interconnected through IoT, this would further exacerbate this data requirement and indexing capability concerns. Tools are needed to alleviate this demand.

## Four Pillars of Observability



**Monitoring**



**Alert / Visualization**



**Distributed Systems  
Tracing Infrastructure**



**Log Aggregation /  
Analytics**

*The 4 pillars of observability from the Twitter  
Observability Engineering team’s charter*

## 4 Debug Faster with Continuous Code Improvement

Given that a considerable amount of developer time is spent on debugging, it would make sense to use tools to automate and consolidate the three strategies described above. A continuous code improvement (CCI) platform provides all three of the solutions discussed above for easier debugging.

- ✓ It automates much of the effort in debugging errors through grouping, deduping, triaging, integration into communication channels for alerts, and more.
- ✓ It aggregates logs from all your stacks, environments, and systems, with an intuitive interface to search and understand those mountains of information.
- ✓ It provides deep observability into your systems through real-time contextual information about errors.

Let's look at a few of these features in detail: grouping, contextual information, and connecting information across environments.

### Grouping

Developers want a solution to let them proactively deal with bugs, but most solutions can't accurately identify unique bugs. Instead, they treat every bug as a unique instance, leaving developers stuck with noisy solutions or learning about errors from customers. They have to be reactive and comb through errors manually, dealing with notification spam or searching through the errors as if they were logs.

A CCI solution uses machine learning to determine error patterns and identify error types to understand when errors are the same or different. A fingerprint is given to identify every occurrence of an event and then combines events accordingly. If it sees the same root cause, for example, it groups those errors into one occurrence. If it sees similar exception class names and properties, it combines them into another occurrence. This method gets rid of missed bugs and noise and makes it easy to not only classify and prioritize errors but also automate the response to errors. Because errors are grouped accurately, a CCI solution can automatically trigger workflows based on any new bugs or regressions that are detected to proactively address issues and minimize their impact.

### Contextual Information

Logs are great. And lots of them are even better, but they are also a huge problem when it comes to debugging. There are so many logs that finding a particular problem is like looking for a needle in a haystack.

On the other end of the spectrum, there is APM. It's great for following flows through an application and between applications. But APM solutions are more focused on latency and less focused on errors.

CCI fills the middle ground, providing all the code-context and contextual metadata needed to move quickly and resolve errors. The stack trace is revealed, along with the exact line of code that caused the error and the related git blame information. HTTP request parameter values, local variable values that happened at runtime, and more, are also exposed.



## Connecting Errors Across All Environments

By using CCI across all environments in the SDLC, users gain observability.

In test environments, CCI accelerates testing by identifying root cause and helping to quickly triage and communicate errors. In staging, CCI improves release readiness by quickly identifying and tracking errors against complex production-mirror environments. And in production, CCI implements live monitoring on production apps so errors in your code are alerted and understood before customers are affected. But best of all, when CCI is used across all these environments, it adds even greater value, by providing a holistic view of errors, including log messages, where the error first occurred, if the error has been fixed in a newer version/ environment, and occurrences of the error in each environment. Referring back to the testing discussion above, CCI makes it easier for teams to make up the time cost by reducing the need to get all the way to p=1 and also cutting the dollar cost at the same time.

CCI provides true insight into why, how, and where errors are occurring at all stages of development.

## 5 Conclusion

Testing is an imperfect process, and debugging can be expensive. However, there are strategies and tools that can help reduce the time and effort in debugging, and increase the quality of code and reliability of deployments.

Continuous code improvement solutions can help track down and fix errors, turn tests from red to green faster, reduce production issues, improve developer confidence, and ultimately improve business bottom lines.



## About Rollbar

Rollbar is the leading continuous code improvement platform that proactively discovers, predicts, and remediates errors with real-time AI-assisted workflows. With Rollbar, developers continually improve their code and constantly innovate rather than spending time monitoring, investigating, and debugging. More than 5,000 businesses, including Twilio, Salesforce, Twitch, and Affirm, use Rollbar to deploy better software, faster while quickly recovering from critical errors as they happen. Learn more at [rollbar.com](https://rollbar.com)



© 2012–21 ROLLBAR, INC.